TrustSketch: Trustworthy Sketch-based Telemetry on Cloud Hosts

Zhuo Cheng, Maria Apostolaki, Alan Liu, Vyas Sekar

Carnegie Mellon University

PRINCETON UNIVERSITY

UNIVERSITY OF MARYLAND

Network telemetry in cloud is important

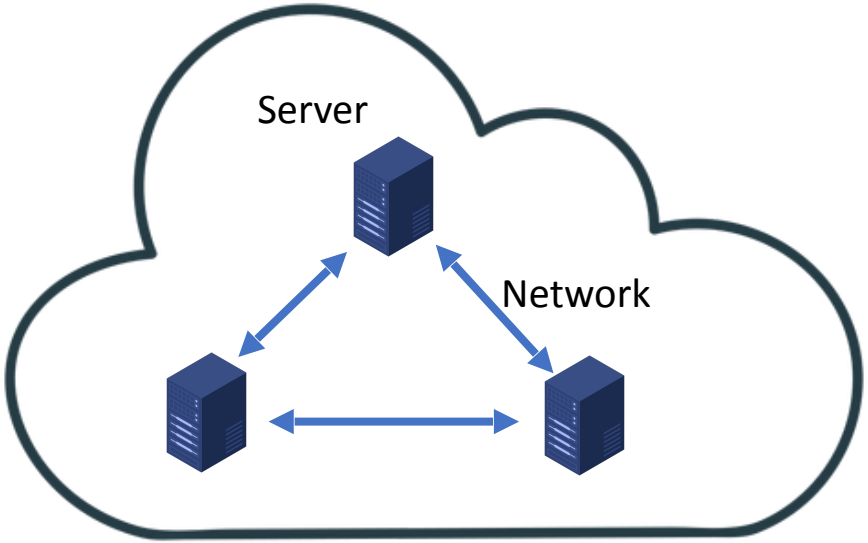

aws

Google Cloud

Microsoft Azure

Network telemetry in cloud is important

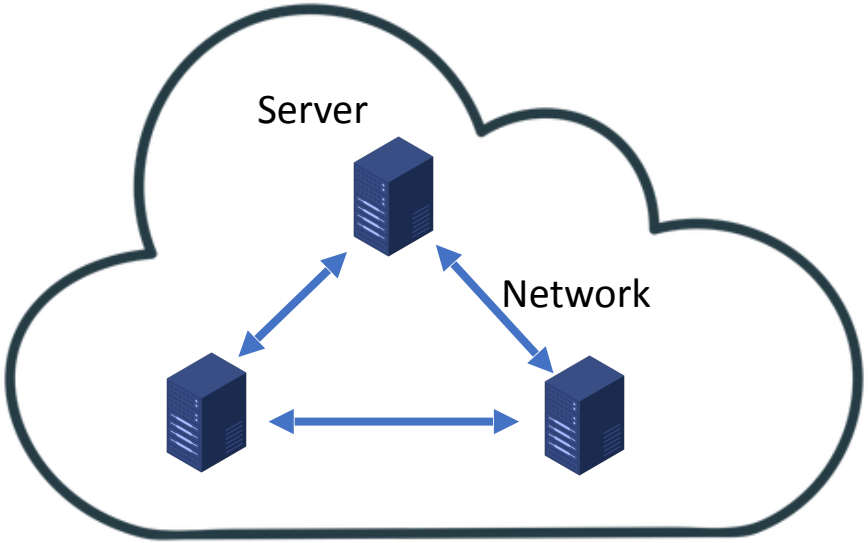

aws

Google Cloud

Microsoft Azure

Network Telemetry

Flow	Volume
IP_A -> IP_B	1GB
IP_A -> IP_C	5GB

Operator

Network telemetry in cloud is important

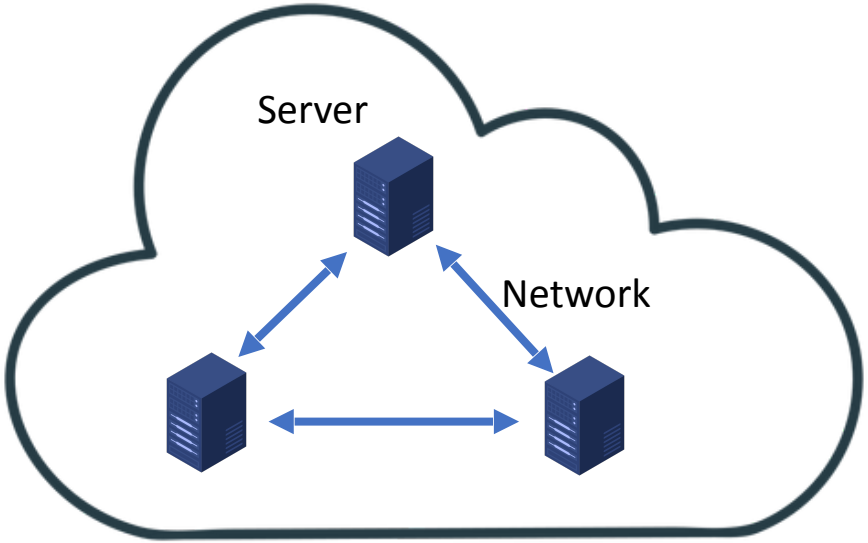

aws

Google Cloud

Microsoft Azure

Network Telemetry

Flow	Volume
IP_A -> IP_B	1GB
IP_A -> IP_C	5GB

Operator

DDoS attack monitoring

Accounting

Background: Sketch-based telemetry

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

Count-min Sketch

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

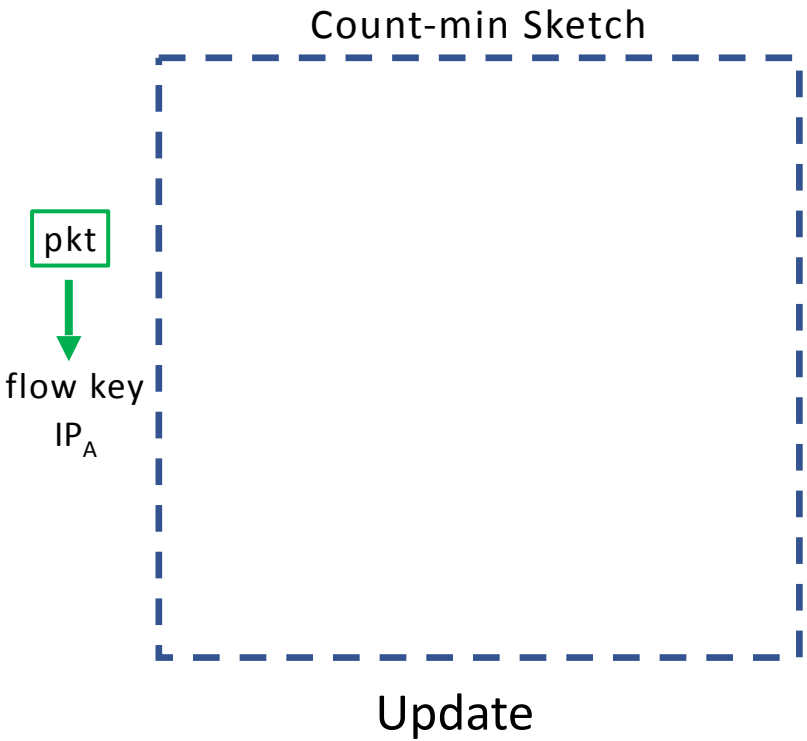

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

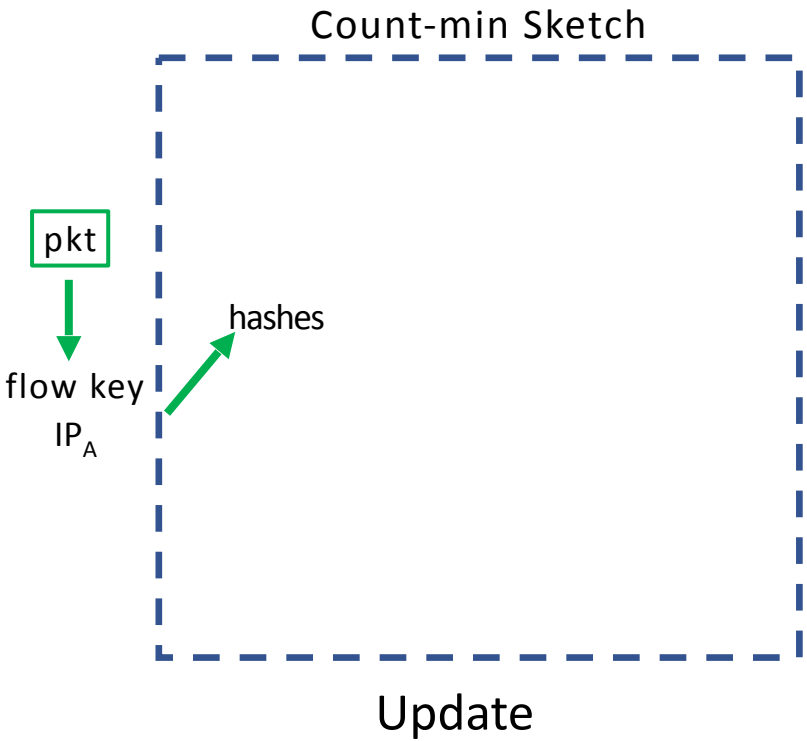

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

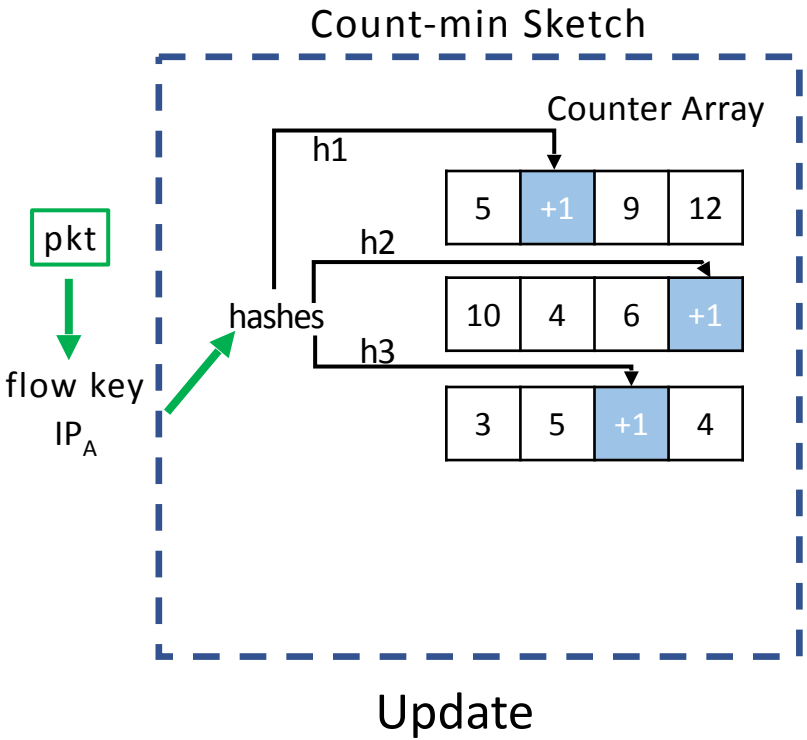

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

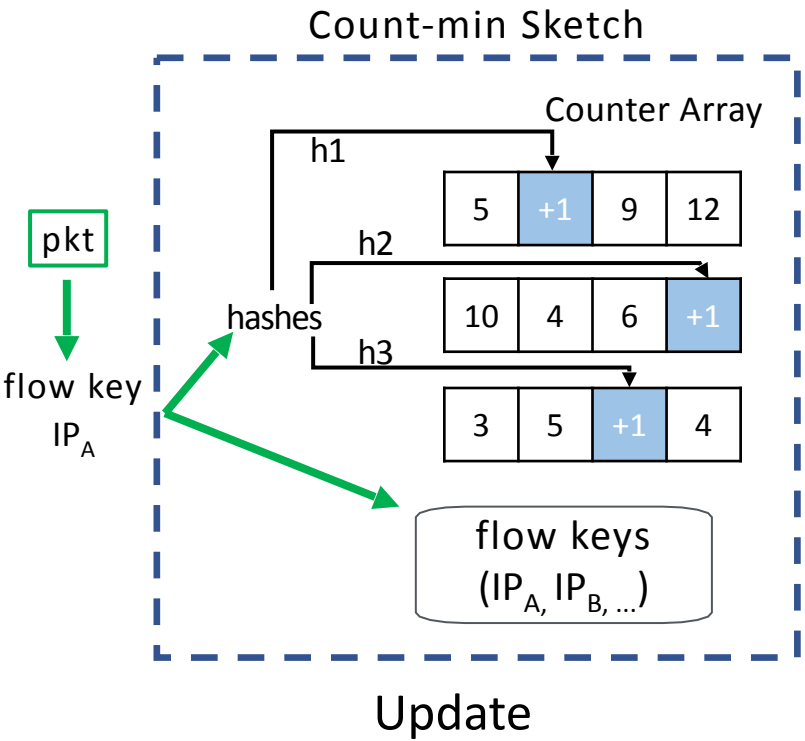

How much data been sent from each IP?

Operator

Background: Sketch-based telemetry

Low footprint, Efficient, Accurate

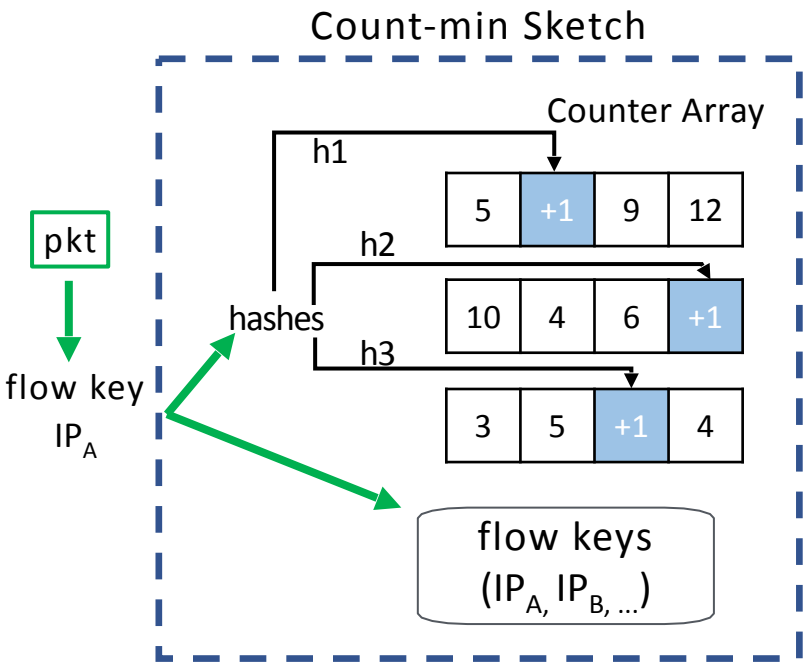

Update

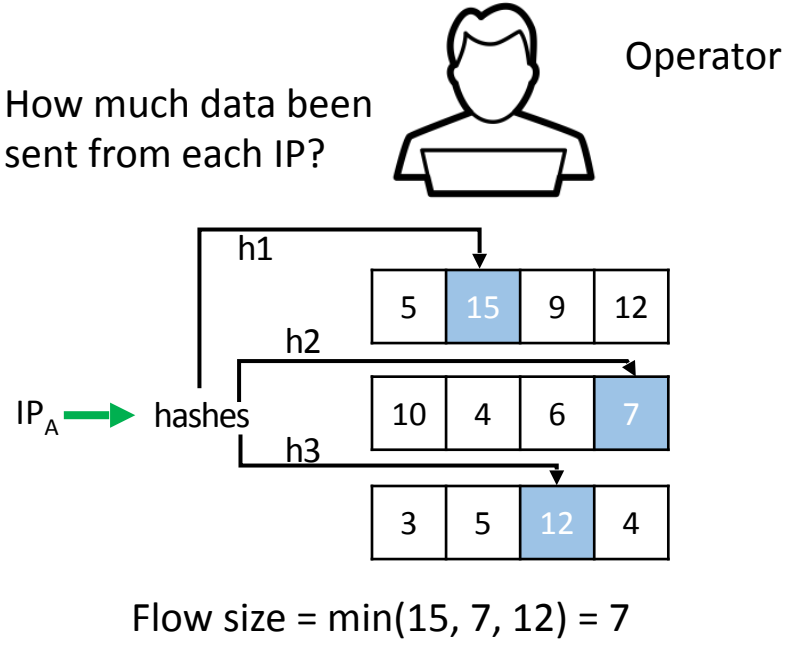

Query

Prior work proposes to run sketches on hosts

UnivmonSketch (SIGCOMM'16), SketchVisor (SIGCOMM'17), ElasticSketch (SIGCOMM'18), NitroSketch (SIGCOMM'19), CocoSketch (SIGCOMM'21), OctoSketch (NSDI'24)

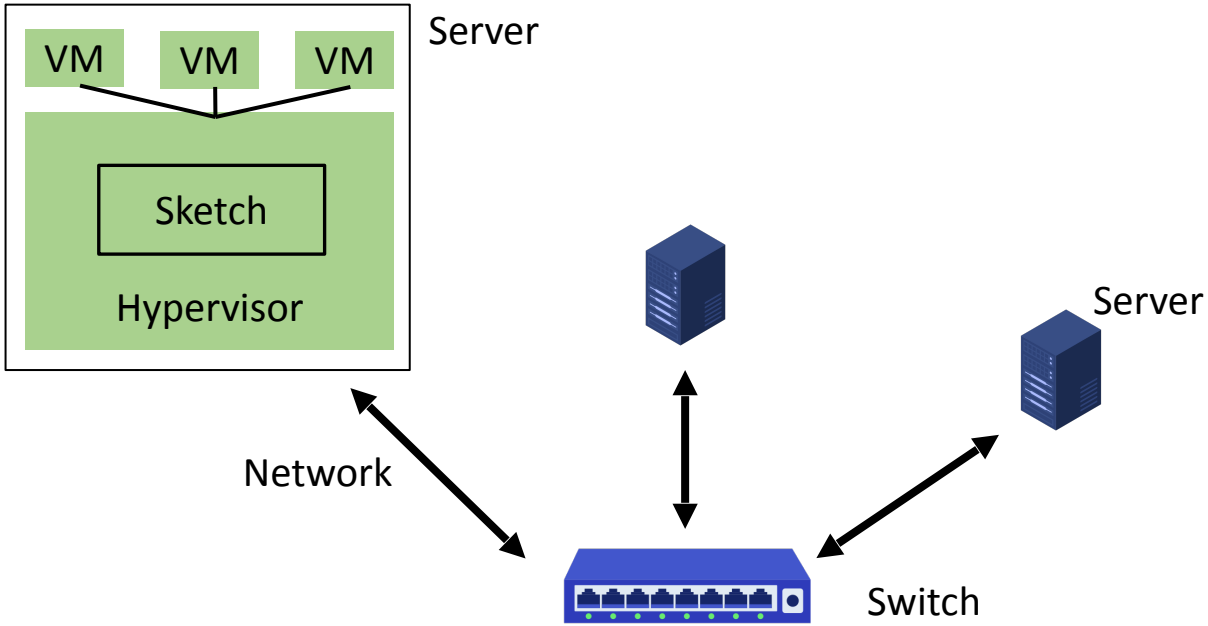

Prior work proposes to run sketches on hosts

UnivmonSketch (SIGCOMM'16), SketchVisor (SIGCOMM'17), ElasticSketch (SIGCOMM'18), NitroSketch (SIGCOMM'19), CocoSketch (SIGCOMM'21), OctoSketch (NSDI'24)

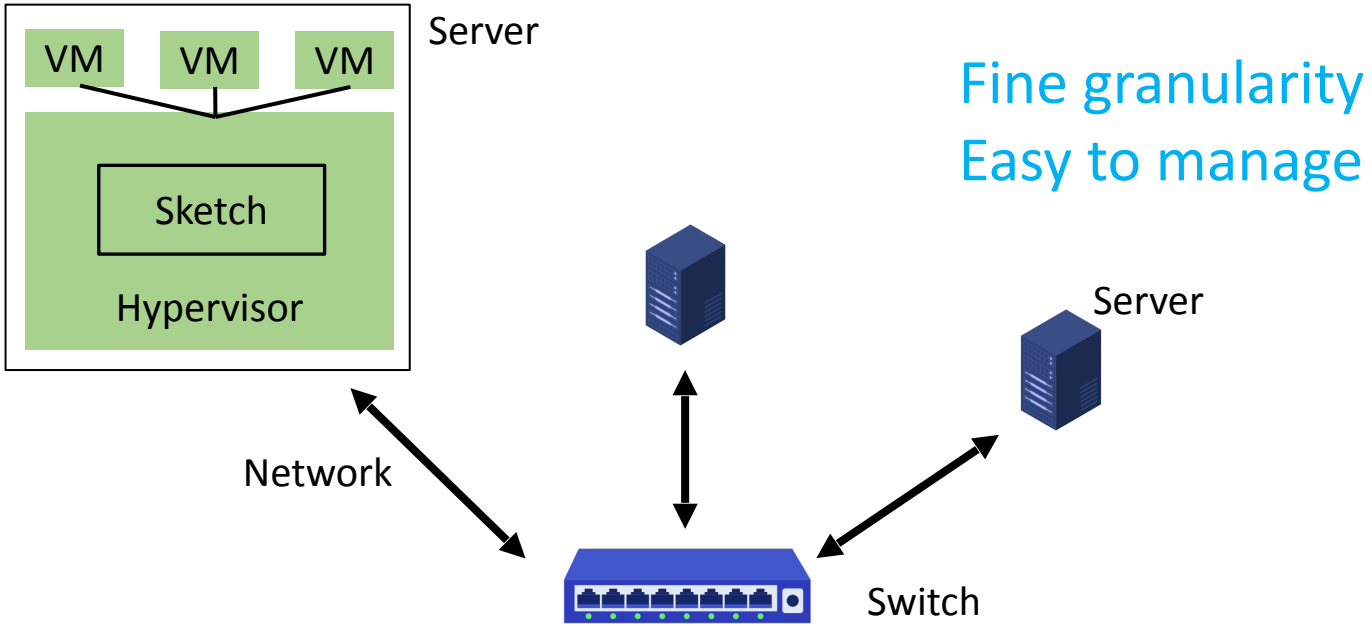

Problem: Sketch compromise impacts downstream task

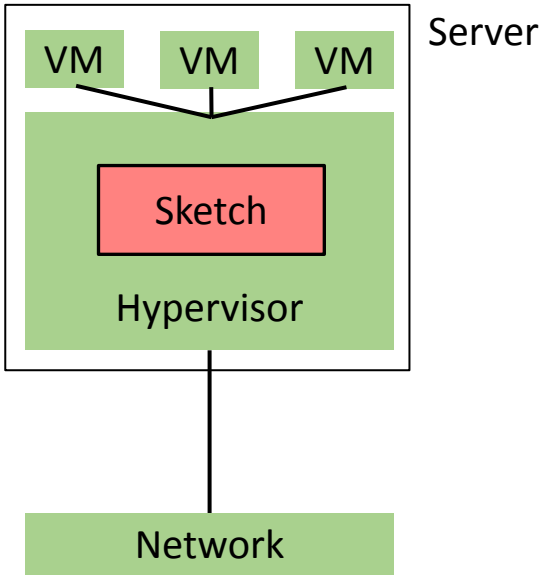

Problem: Sketch compromise impacts downstream task

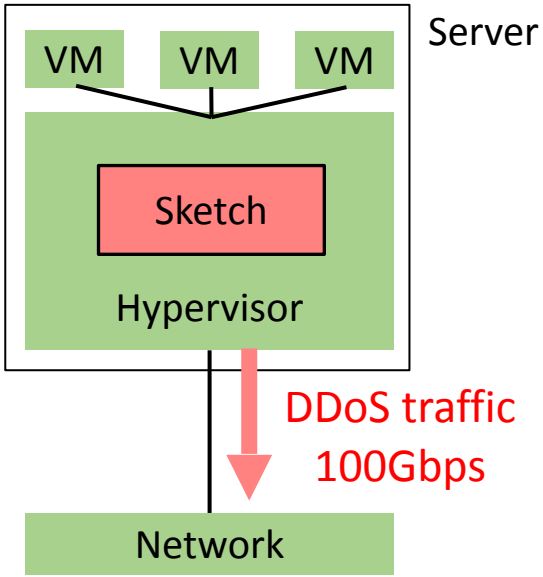

Problem: Sketch compromise impacts downstream task

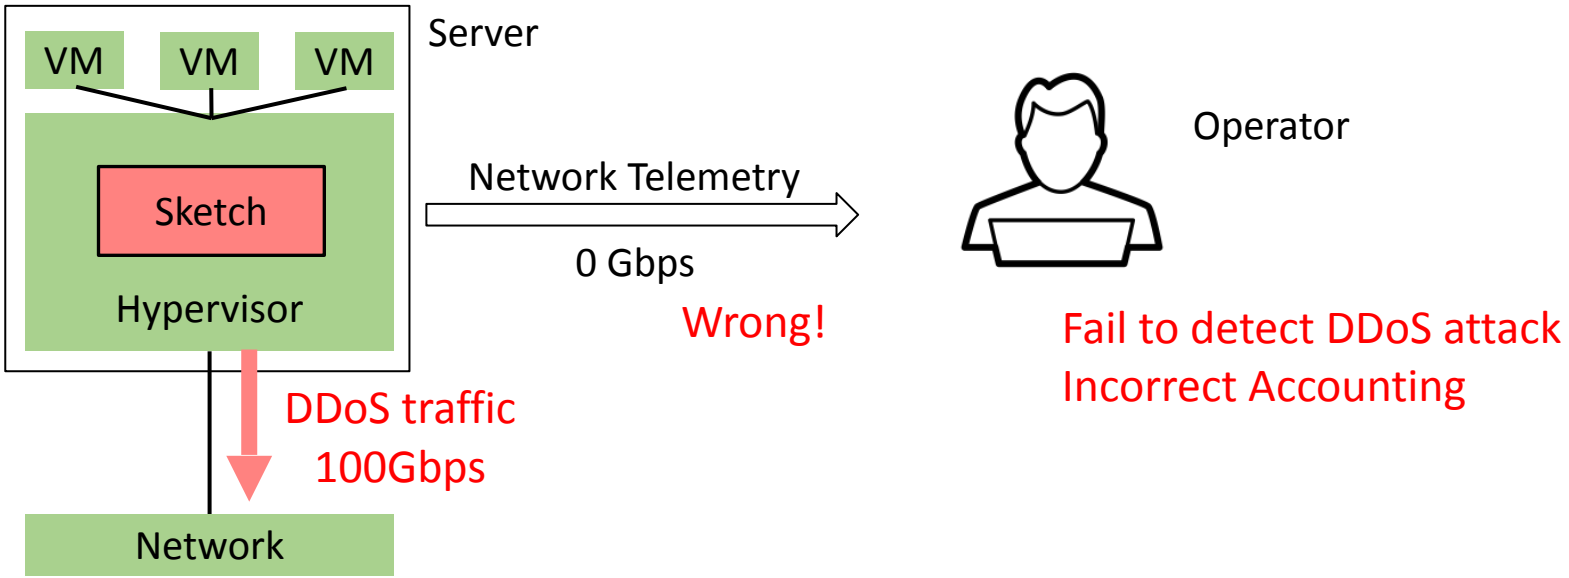

Problem: Sketch compromise impacts downstream task

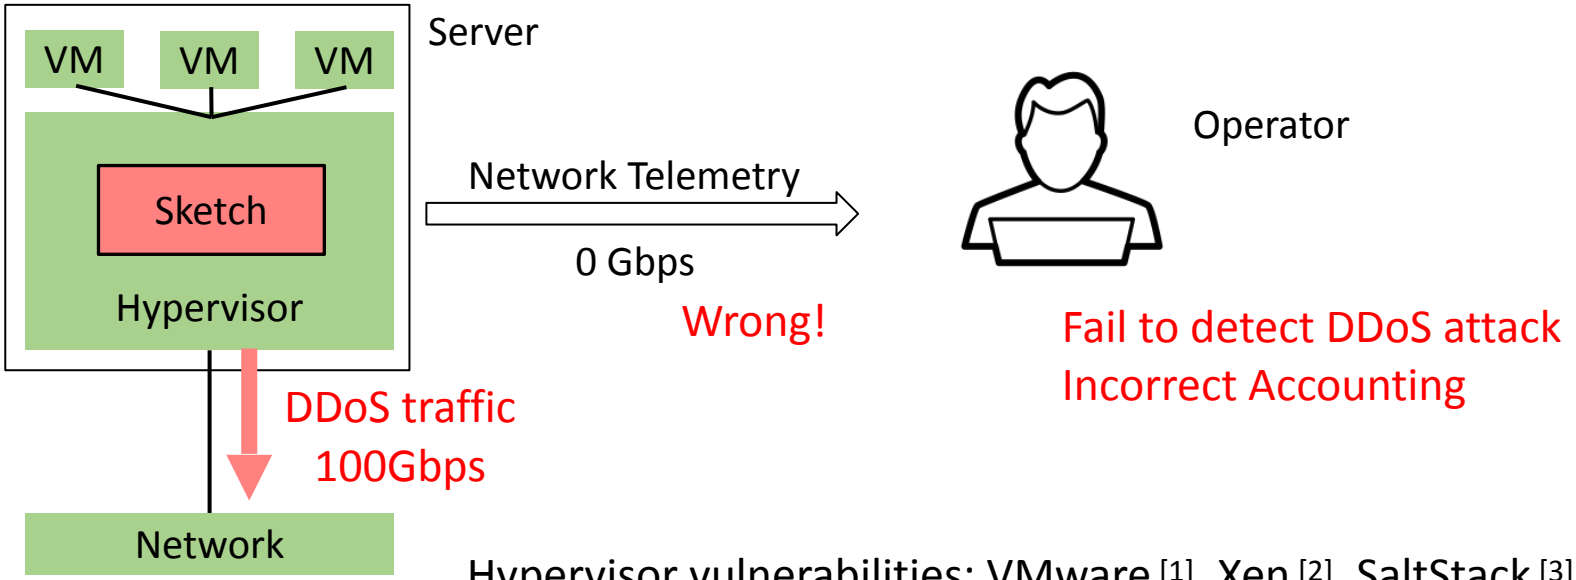

Hypervisor vulnerabilities: VMware [1], Xen [2], SaltStack [3]

[1] https://www.virtualizationhowto.com/2018/11/vmware-esxi-successful-vm-escape-at-geekpwn2018-security-patch/
[2] https://thenewstack.io/privilege-escalation-information-leak-flaws-patched-xen-hypervisor/
[3] https://www.datacenterknowledge.com/security/hackers-exploiting-saltstack-vulnerability-hit-data-centers

Our goal: Trustworthy Sketch-Based Telemetry

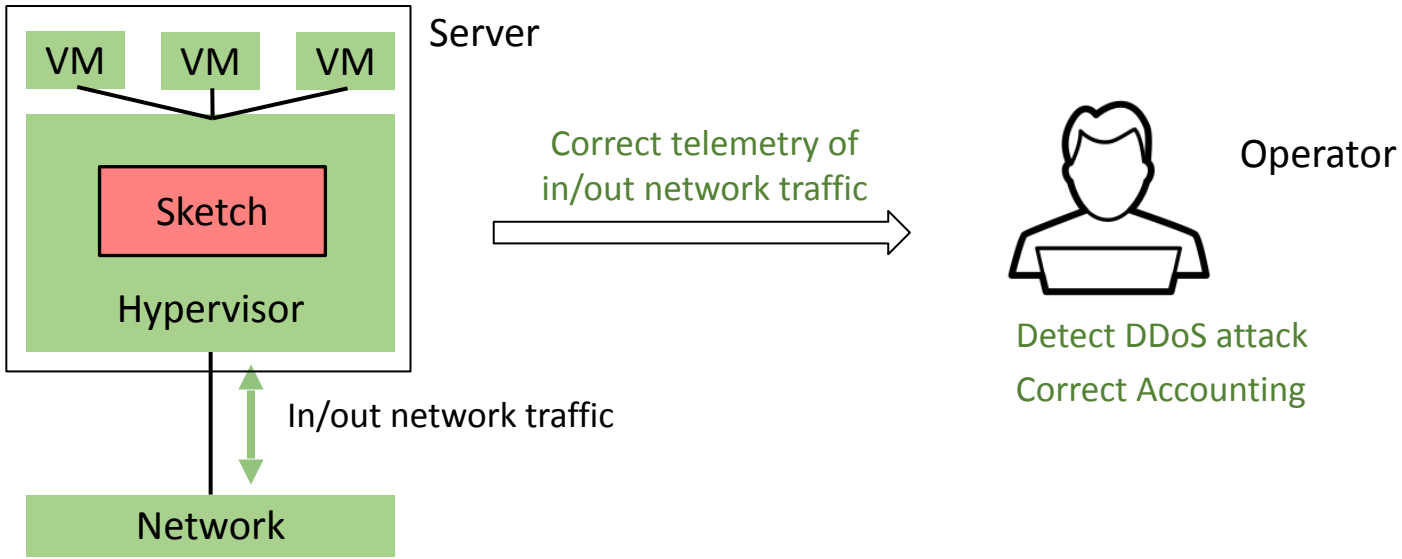

Our contribution

Our contribution

- Formulate requirements for trustworthy sketch-based telemetry.

Our contribution

- Formulate requirements for trustworthy sketch-based telemetry.
- TrustSketch: based on enclave and SmartNIC.

Our contribution

- Formulate requirements for trustworthy sketch-based telemetry.
- TrustSketch: based on enclave and SmartNIC.
- Evaluation shows that TrustSketch is safe with low performance overhead.

Talk Outline

- Motivation.
- Formulate requirements for trustworthy sketch-based telemetry.
- TrustSketch Design.
- Evaluation.

Attacks to compromise the sketch

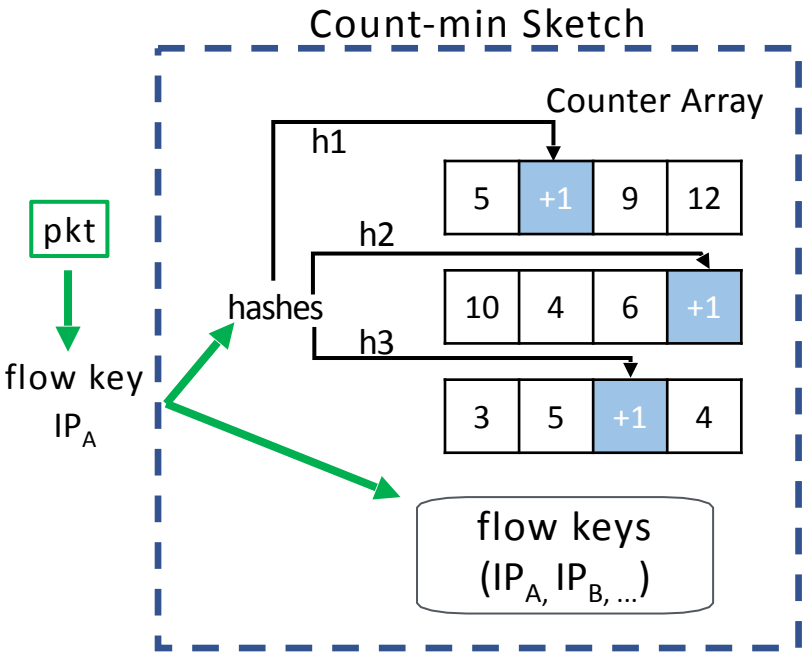

DDoS traffic 100Gbps

IP_A -> IP_{victim}

Attacks to compromise the sketch

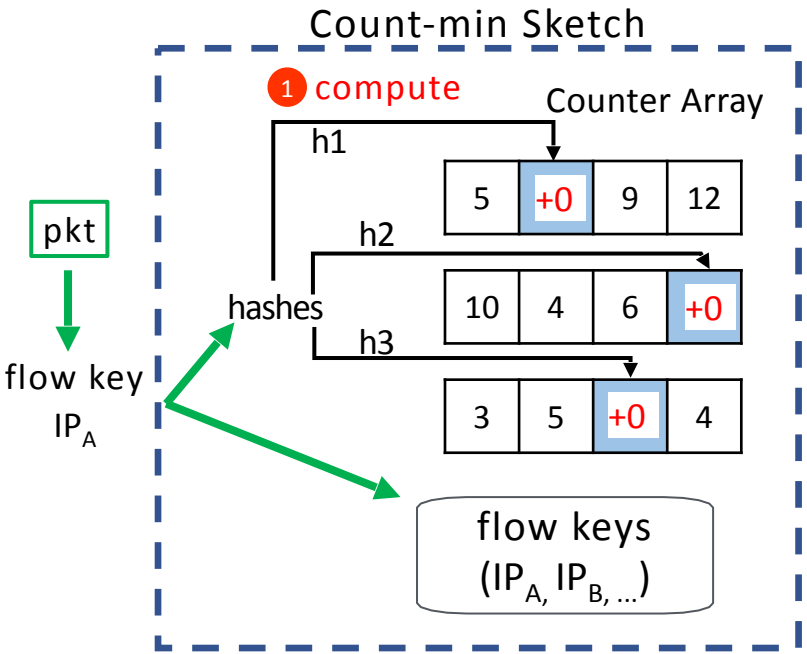

DDoS traffic 100Gbps

IP_A -> IP_{victim}

Attacks to compromise the sketch

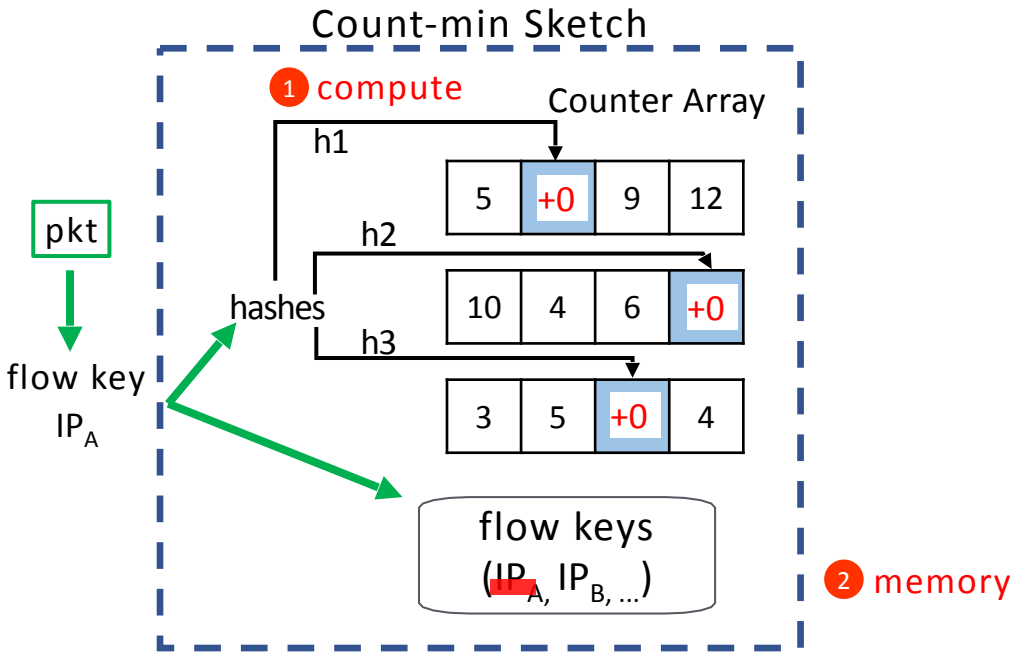

DDoS traffic 100Gbps

IP_A -> IP_{victim}

Attacks to compromise the sketch

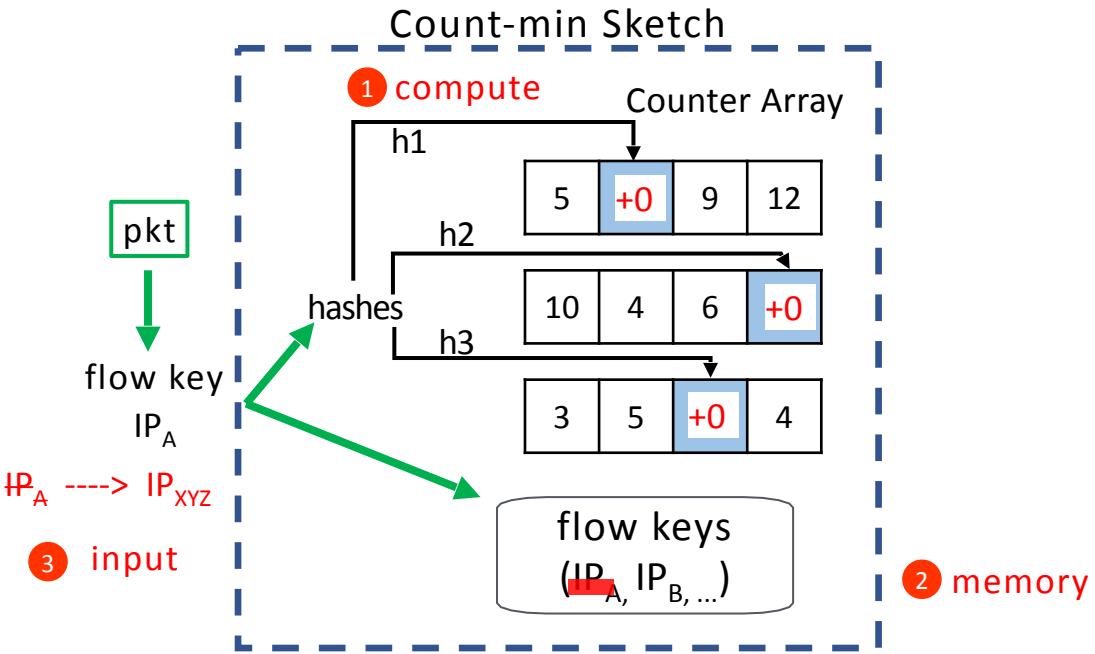

DDoS traffic 100Gbps

IP_A -> IP_{victim}

Requirements for trustworthy sketch-based telemetry

Integrity

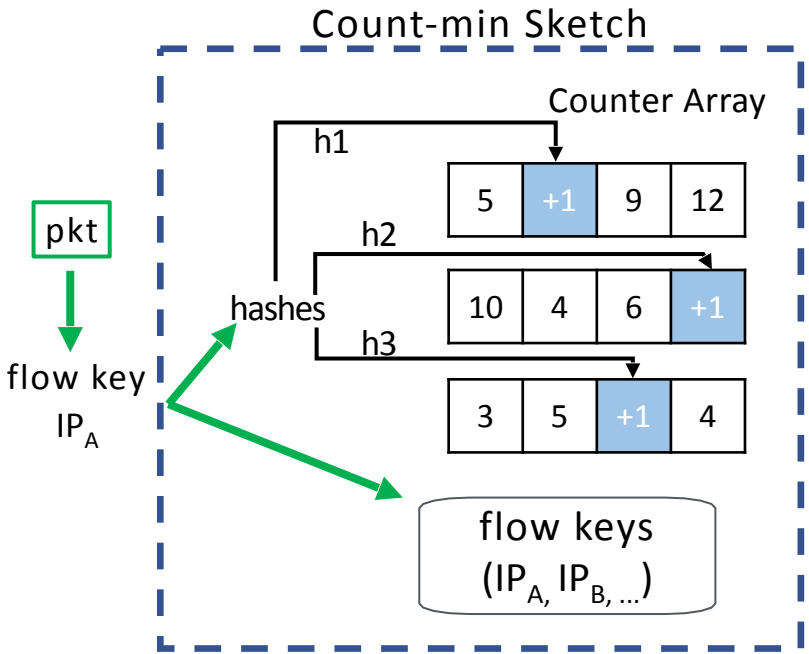

Requirements for trustworthy sketch-based telemetry

Integrity

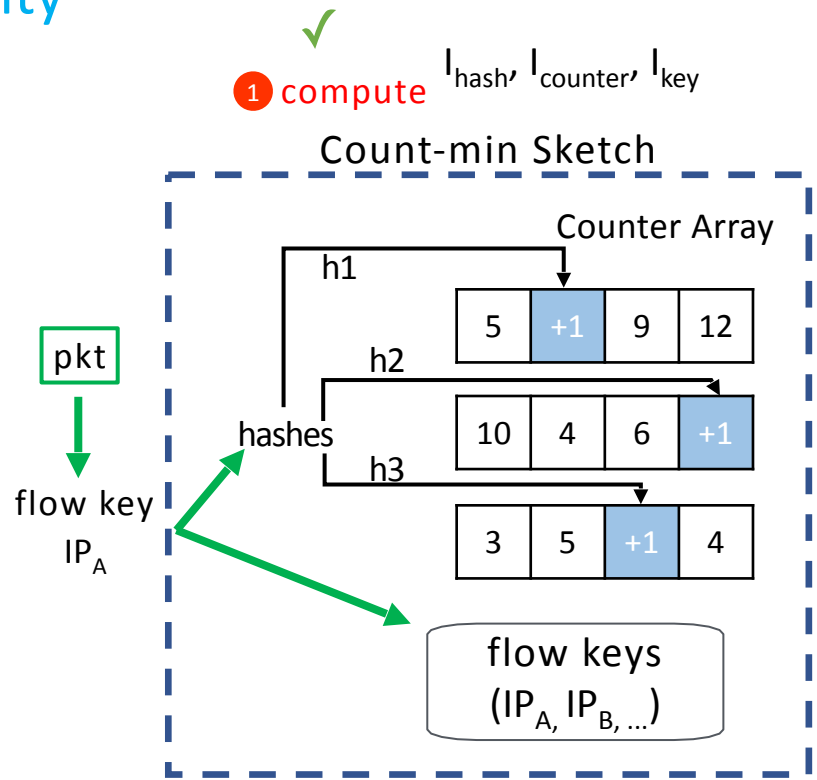

Requirements for trustworthy sketch-based telemetry

Integrity

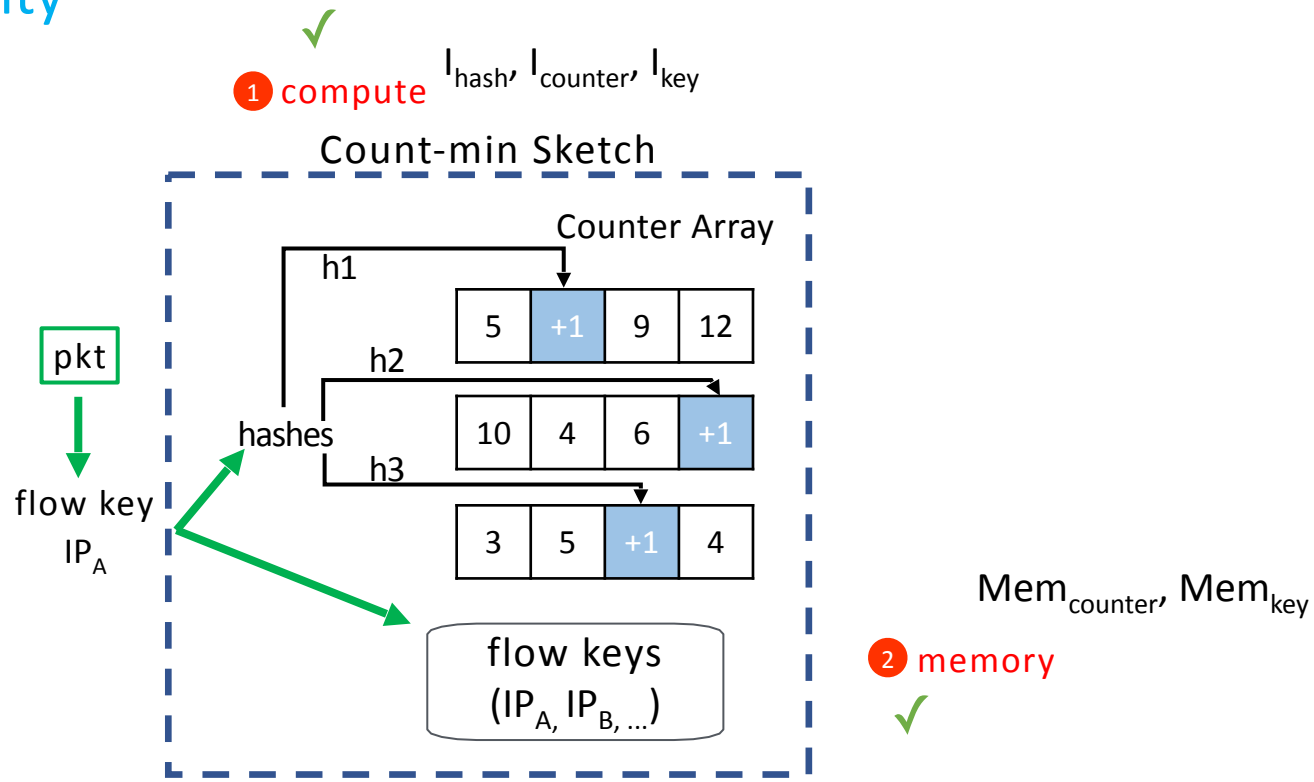

Requirements for trustworthy sketch-based telemetry

Integrity

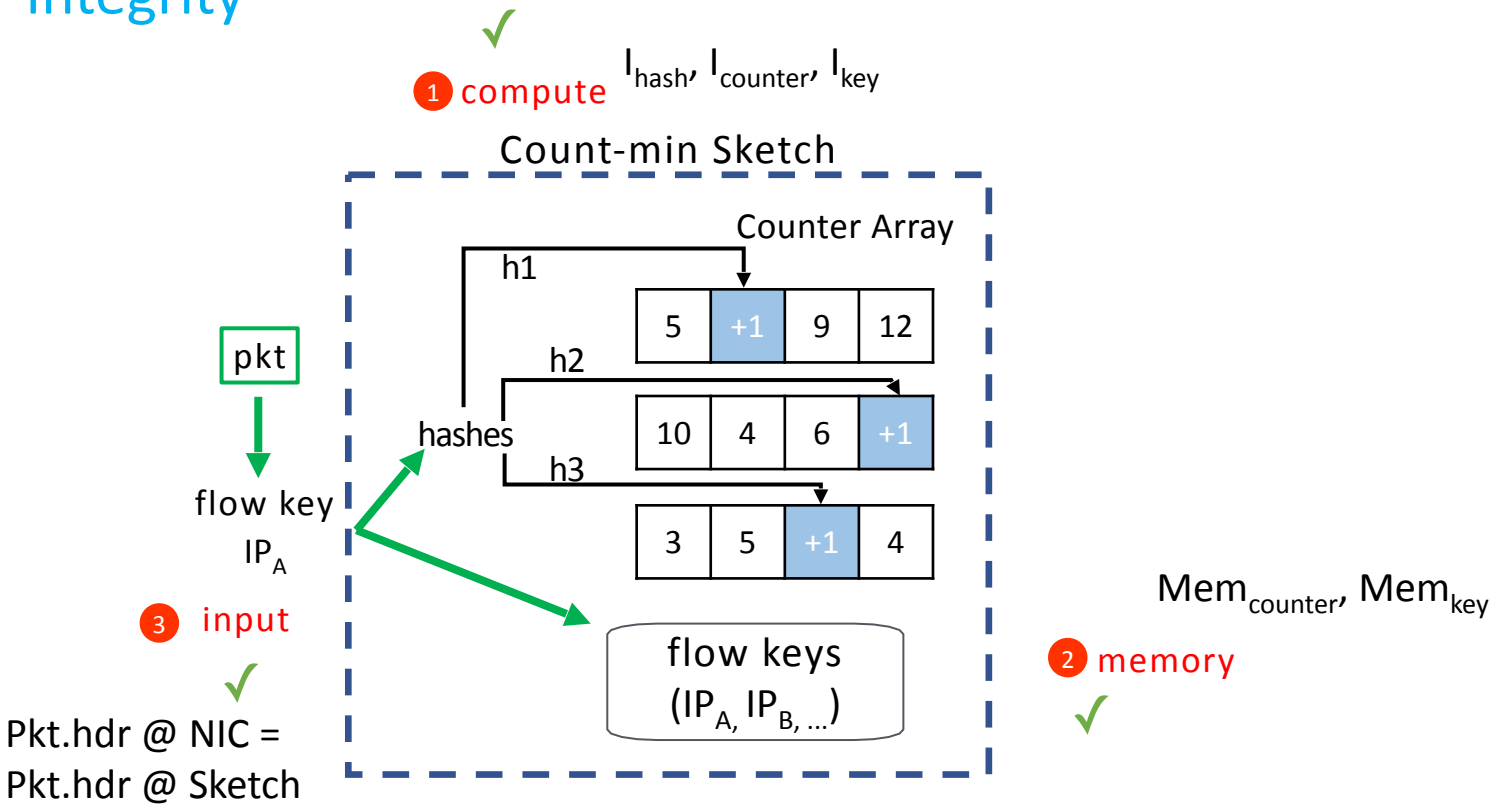

Requirements for trustworthy sketch-based telemetry

Integrity

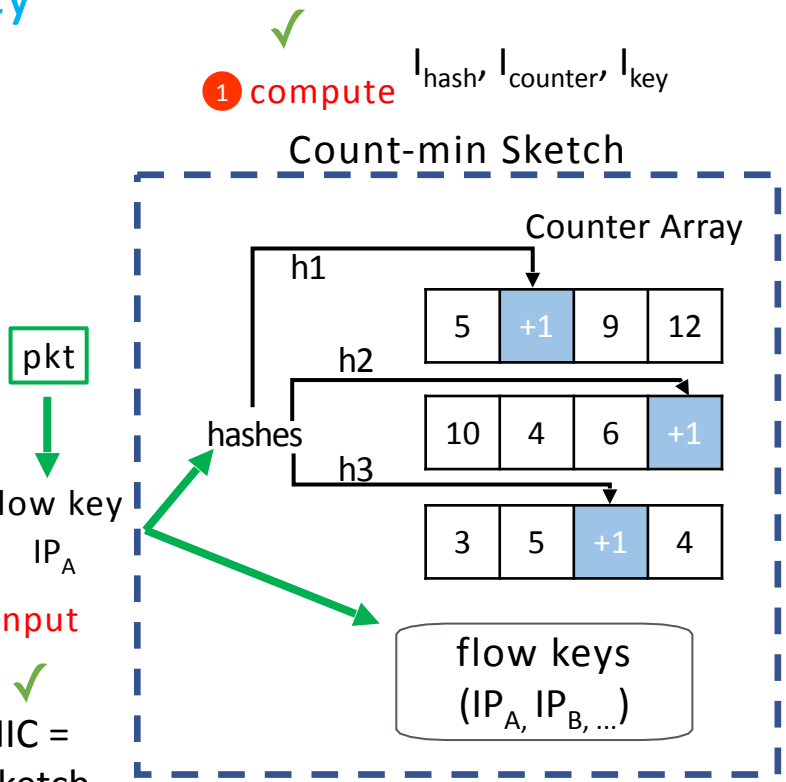

Performance

line-rate

$Mem_{counter}$, Mem_{key}

2 memory

✓

Requirements for trustworthy sketch-based telemetry

Integrity

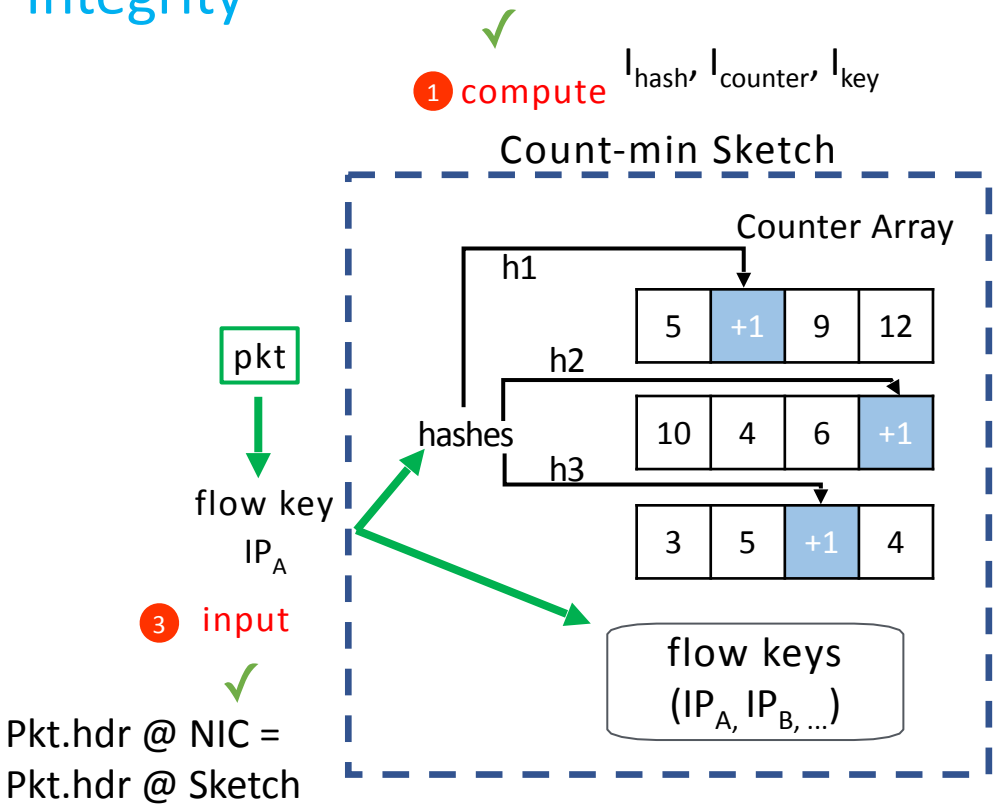

Performance

line-rate

Generality

Sketch 1 -> Flow counting
Sketch 2 -> Heavy hitters
.....

$Mem_{counter}$, Mem_{key}

2 memory

✓

Existing solutions cannot meet the requirements.

Existing Solutions	Integrity			Performance	Generality
	Compute	Memory	Input		
Cross checking[1]					
Code attestation[2]					
Secure memory[3]					

[1] Planck: Millisecond-scale monitoring and control for commodity networks. SIGCOMM 14.
[2] Flexible OS support and applications for trusted computing. HotOS 03.
[3] AMD Secure Memory Encryption (SME).

Existing solutions cannot meet the requirements.

Existing Solutions	Integrity			Performance	Generality
	Compute	Memory	Input		
Cross checking[1]	✓	✓	✓	✗	✓
Code attestation[2]	✗	✗	✗	✓	✓
Secure memory[3]	✗	✓	✗	✓	✓

[1] Planck: Millisecond-scale monitoring and control for commodity networks. SIGCOMM 14.
[2] Flexible OS support and applications for trusted computing. HotOS 03.
[3] AMD Secure Memory Encryption (SME).

Talk Outline

- Motivation.
- Formulate requirements for trustworthy sketch-based telemetry.
- TrustSketch Design.
- Evaluation.

Root of trust: Enclave & SmartNIC

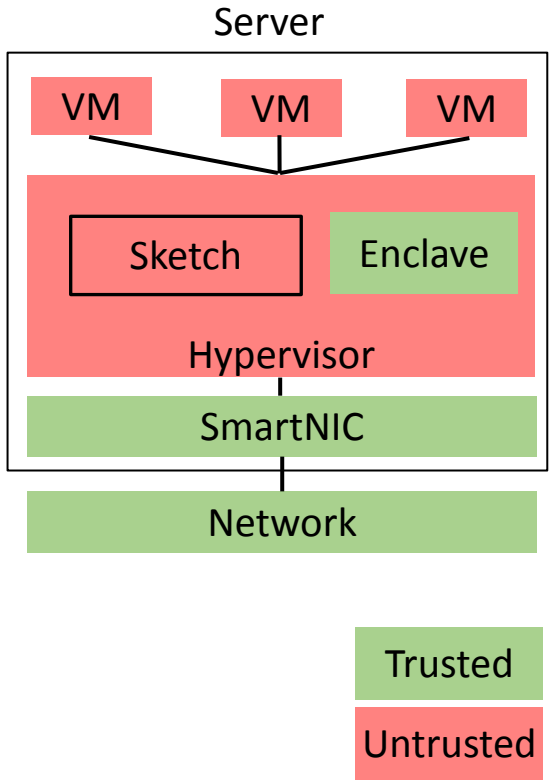

Root of trust: Enclave & SmartNIC

Opportunities:

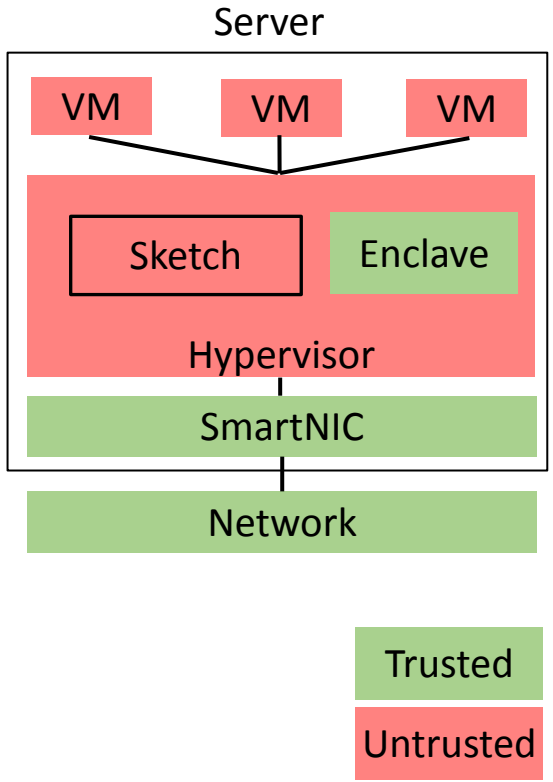

Root of trust: Enclave & SmartNIC

Opportunities:

+ Runtime protection by hardware

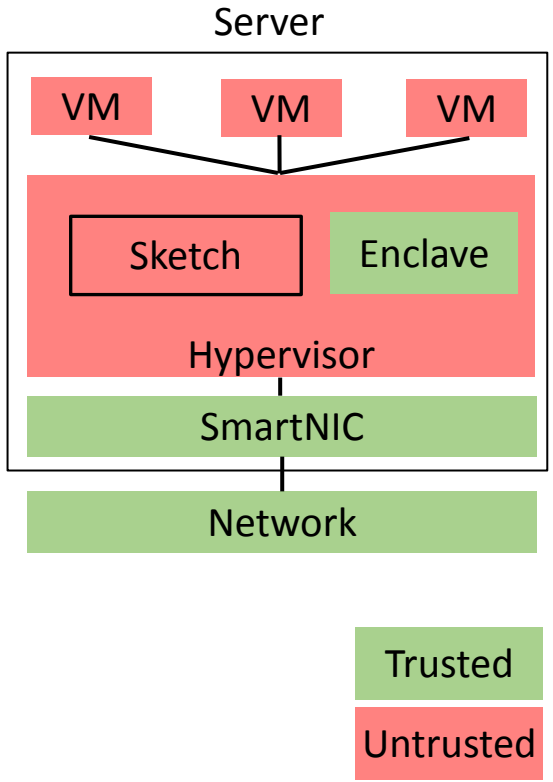

Root of trust: Enclave & SmartNIC

Opportunities:

+ Runtime protection by hardware

+ Already deployed

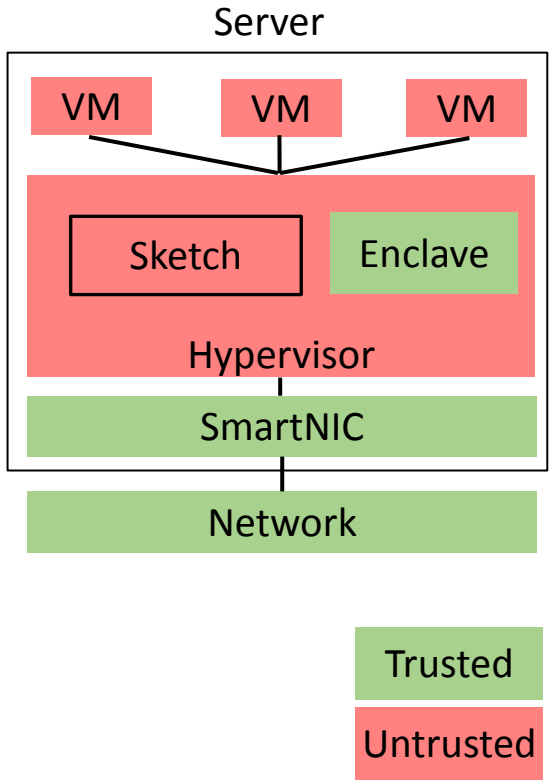

Root of trust: Enclave & SmartNIC

Opportunities:

+ Runtime protection by hardware

+ Already deployed

Constraints:

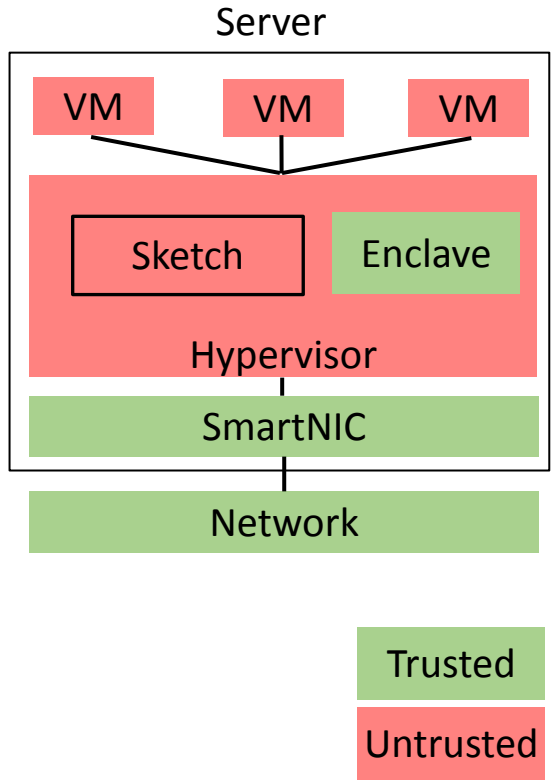

Root of trust: Enclave & SmartNIC

Opportunities:

+ Runtime protection by hardware

+ Already deployed

Constraints:

– Limited compute/memory

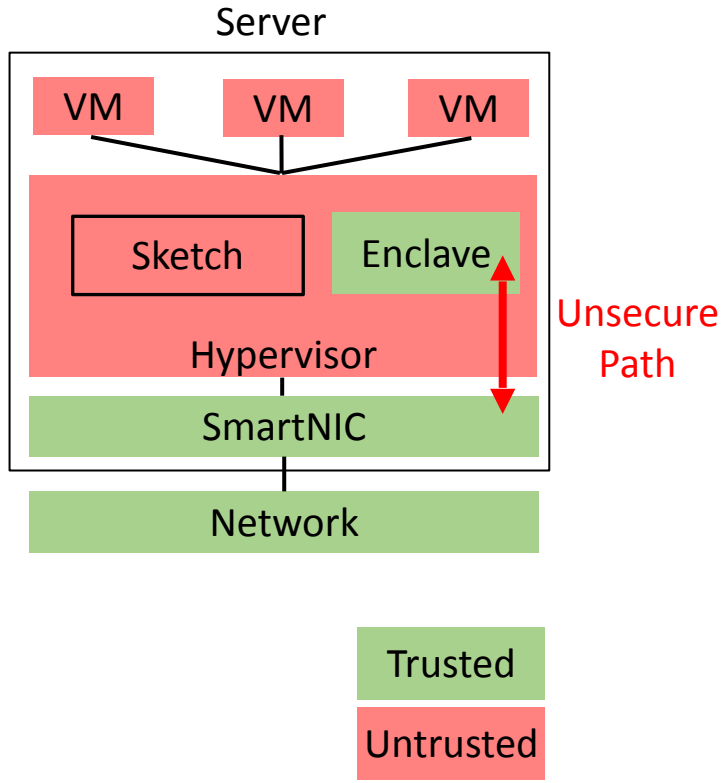

Root of trust: Enclave & SmartNIC

Opportunities:

+ Runtime protection by hardware
+ Already deployed

Constraints:

– Limited compute/memory
– Unsecure path between enclave/NIC

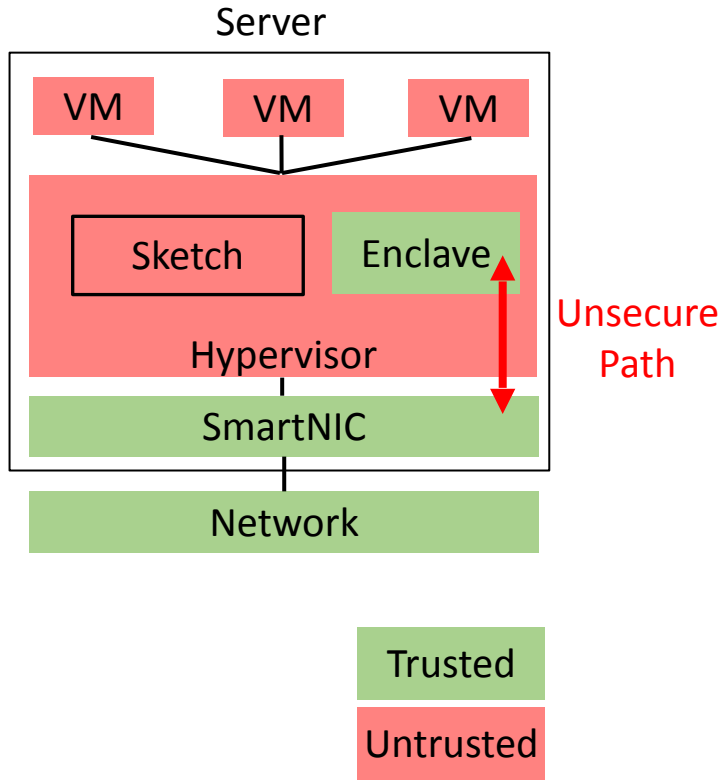

Where to run the telemetry system

Where to run the telemetry system

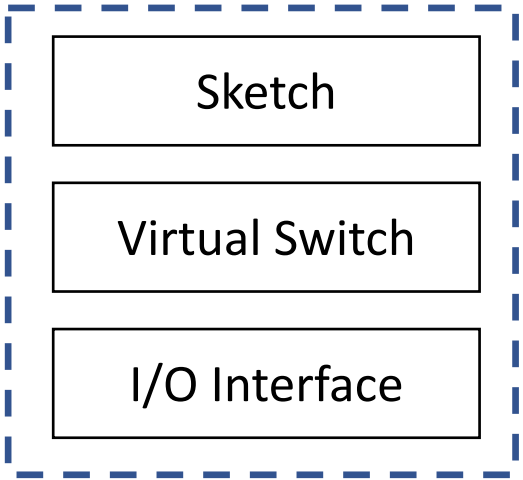

Where to run the telemetry system

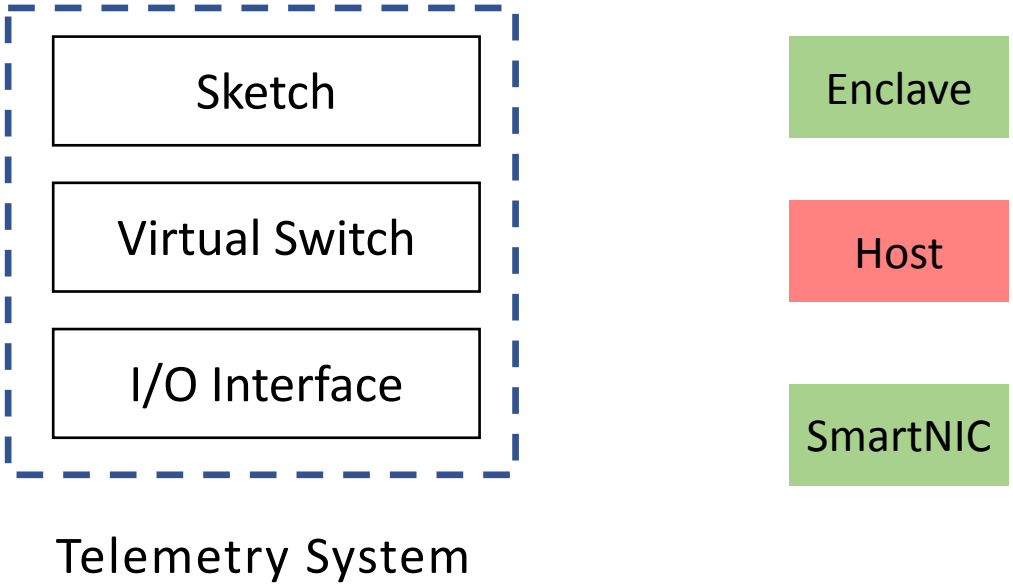

Where to run the telemetry system

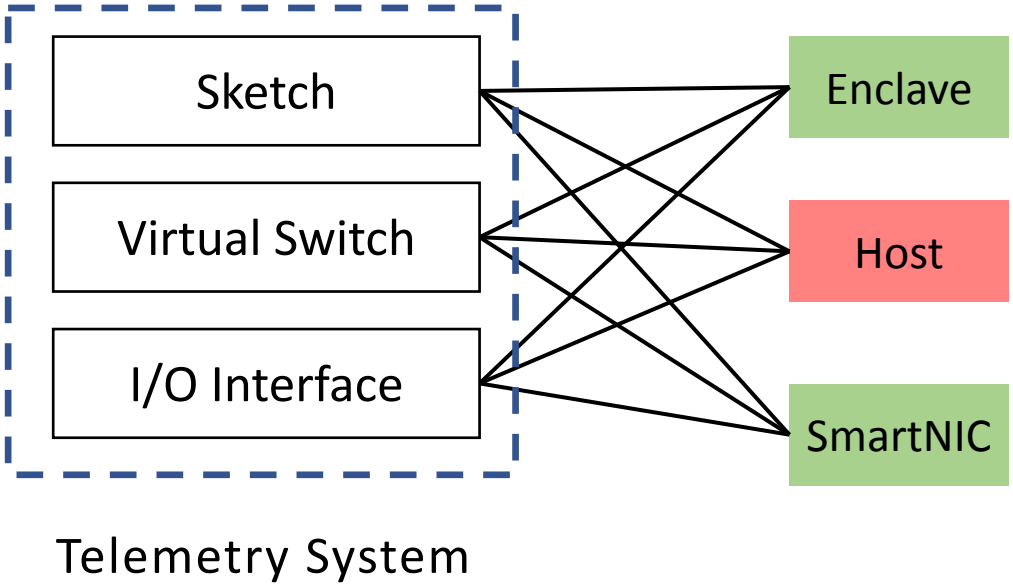

Approach: Smart division of sketch telemetry functions

Approach: Smart division of sketch telemetry functions

Requirements:

1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

Approach: Smart division of sketch telemetry functions

Strawman 1
Sketch in NIC

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

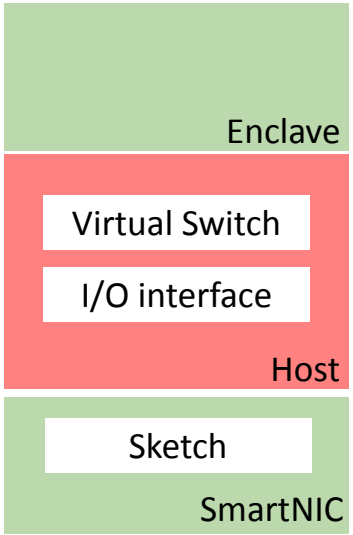

Approach: Smart division of sketch telemetry functions

Strawman 1
Sketch in NIC

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

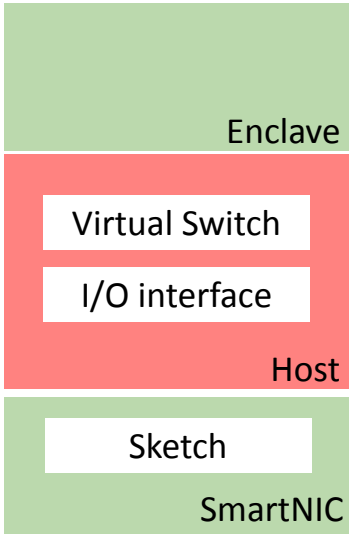

✕ Generality

Approach: Smart division of sketch telemetry functions

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

Strawman 1
Sketch in NIC

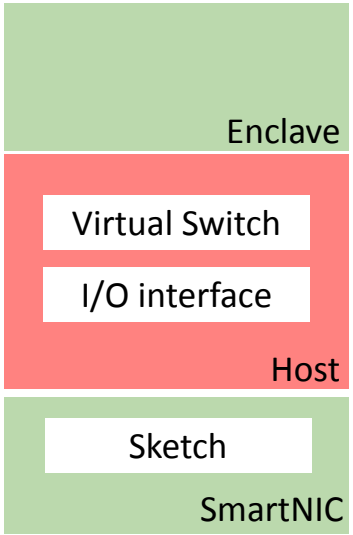

✕ Generality

Strawman 2
All in Enclave

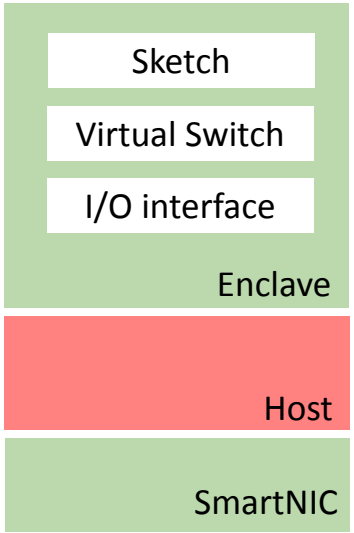

Approach: Smart division of sketch telemetry functions

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

Strawman 1
Sketch in NIC

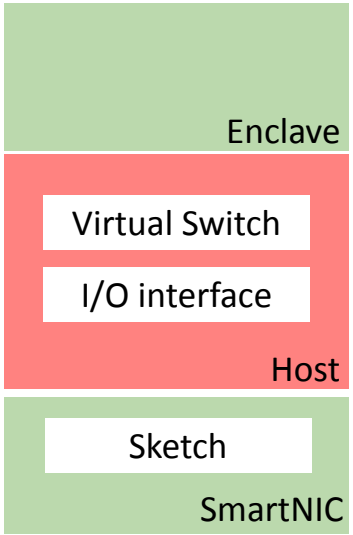

✕ Generality

Strawman 2
All in Enclave

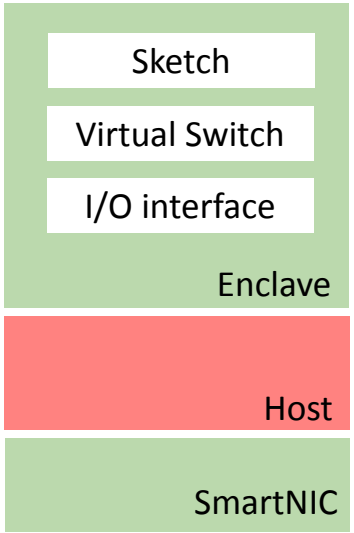

✕ Performance

Approach: Smart division of sketch telemetry functions

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

Strawman 1
Sketch in NIC

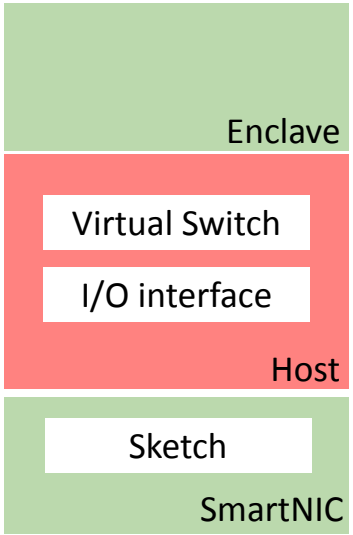

✕ Generality

Strawman 2
All in Enclave

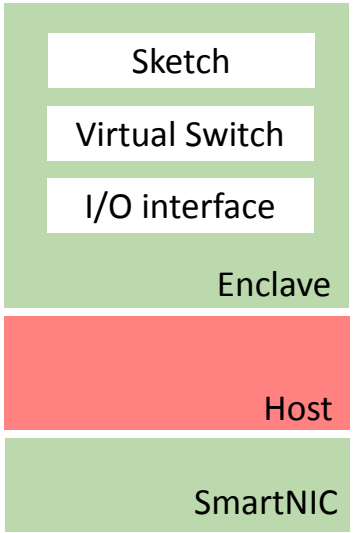

✕ Performance

TrustSketch
Only Sketch in Enclave

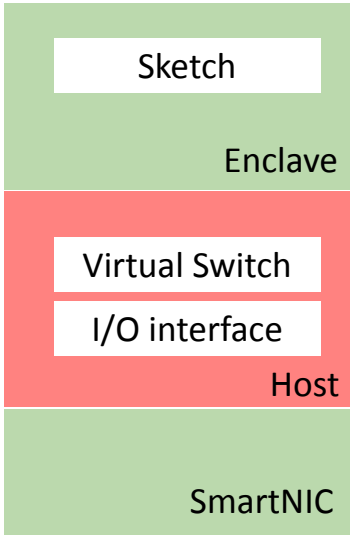

Approach: Smart division of sketch telemetry functions

Requirements:
1. Compute Integrity
2. Memory Integrity
3. Input Integrity
4. Performance
5. Generality

Strawman 1
Sketch in NIC

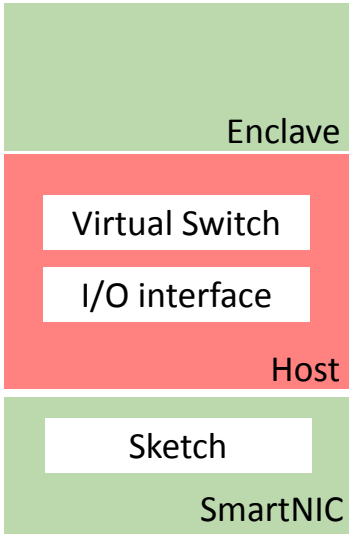

✕ Generality

Strawman 2
All in Enclave

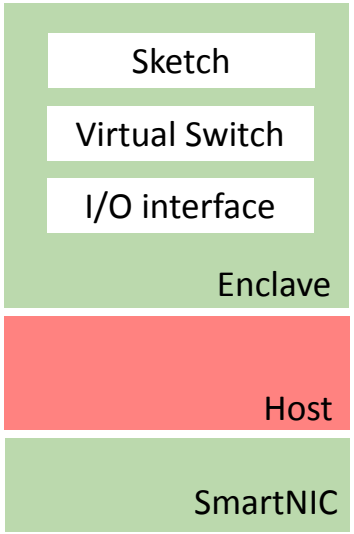

✕ Performance

TrustSketch
Only Sketch in Enclave

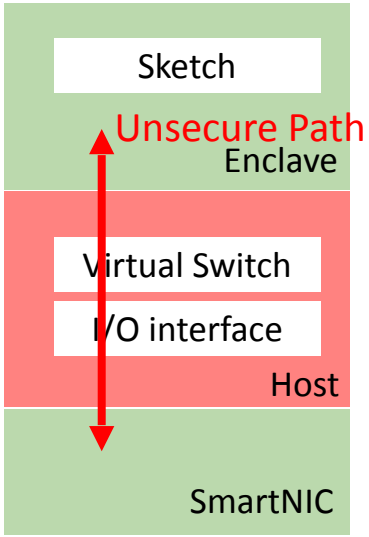

✕ Input Integrity

Leverage SmartNIC to ensure input-integrity

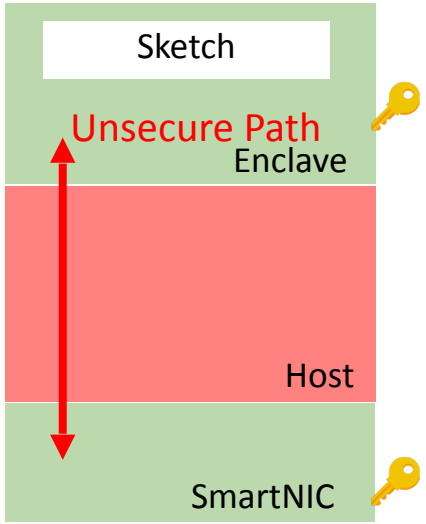

Leverage SmartNIC to ensure input-integrity

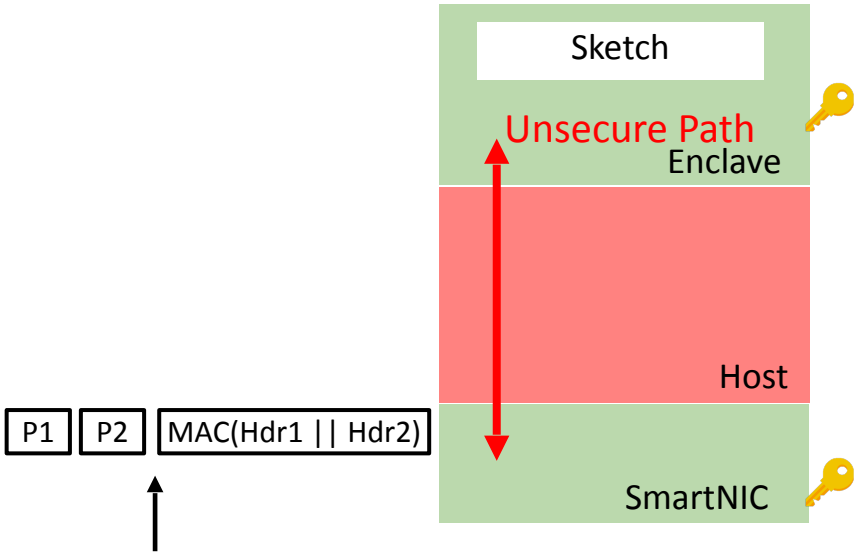

Leverage SmartNIC to ensure input-integrity

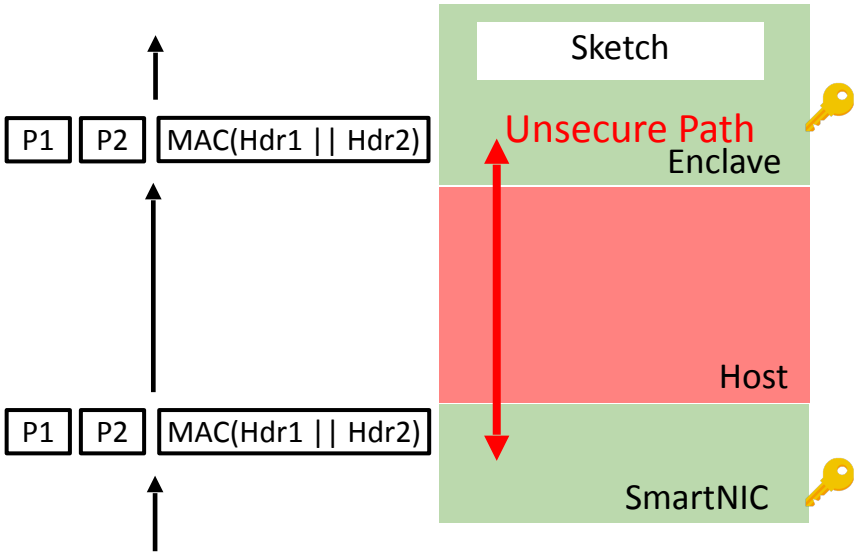

Leverage SmartNIC to ensure input-integrity

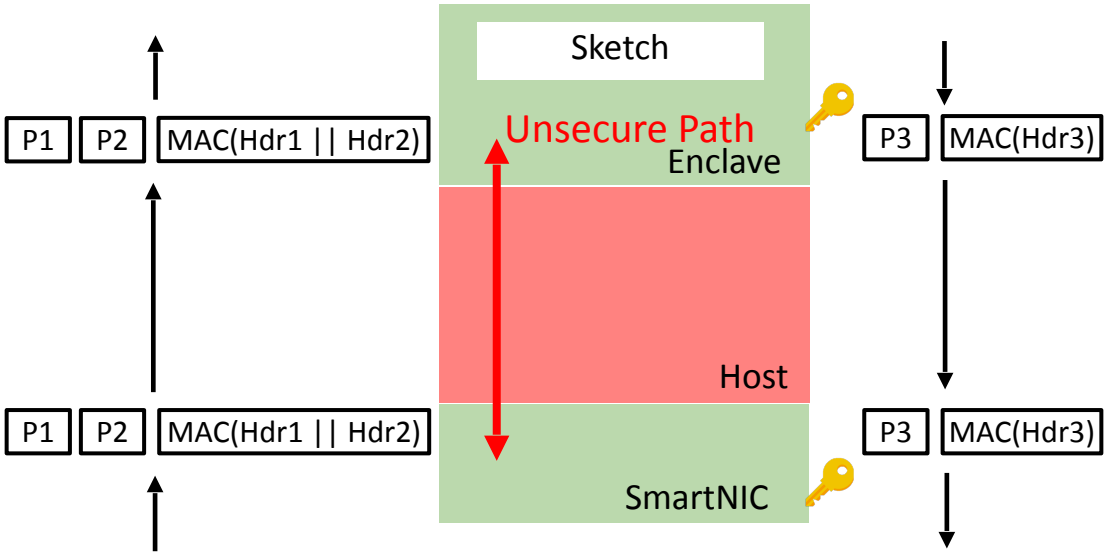

Leverage SmartNIC to ensure input-integrity

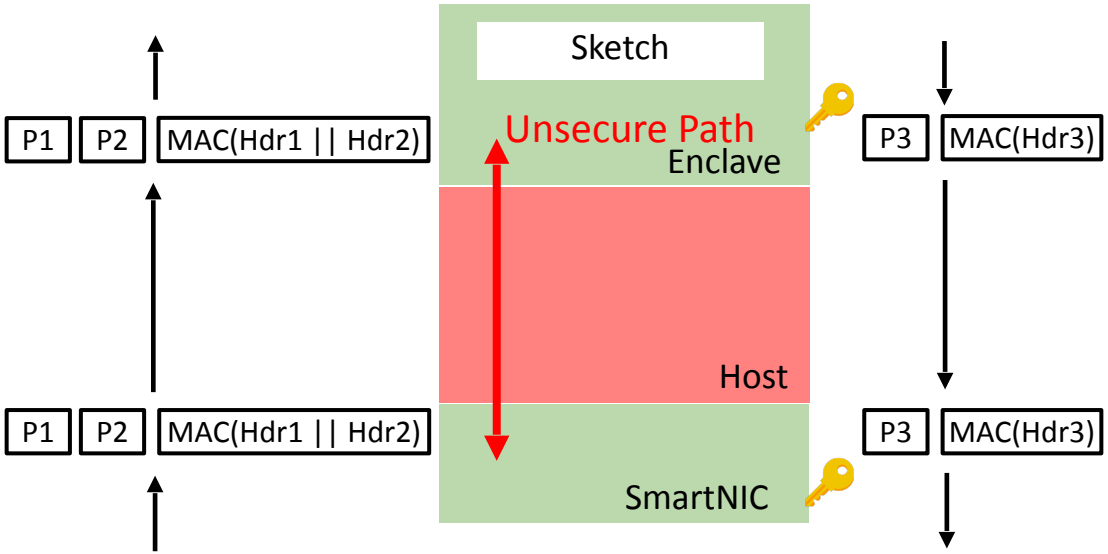

- No packet loss/reorder without attack

Leverage SmartNIC to ensure input-integrity

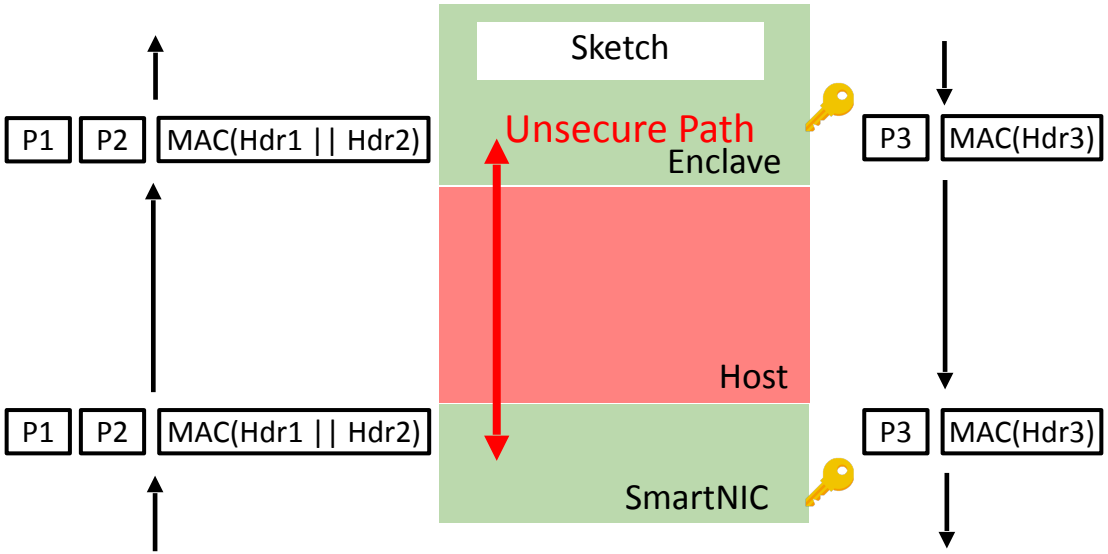

- No packet loss/reorder without attack
- Separate MAC for in/out traffic

Challenge: Handling multi-threaded NICs

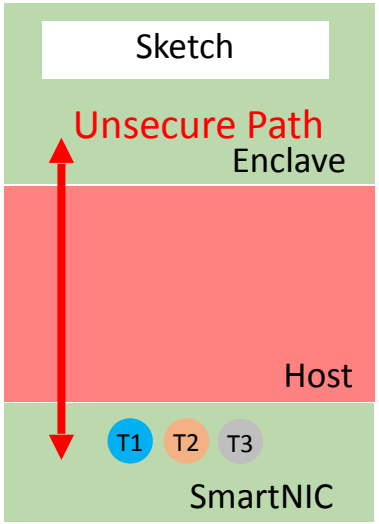

Challenge: Handling multi-threaded NICs

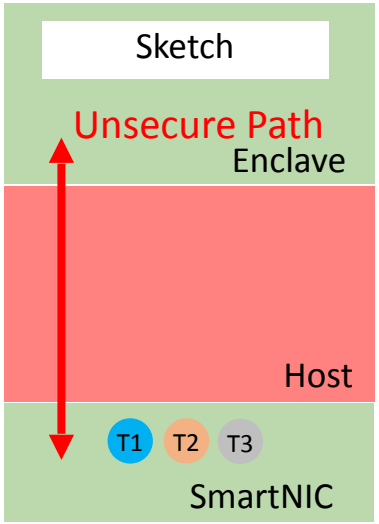

Packet Stream

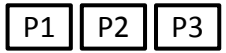

Challenge: Handling multi-threaded NICs

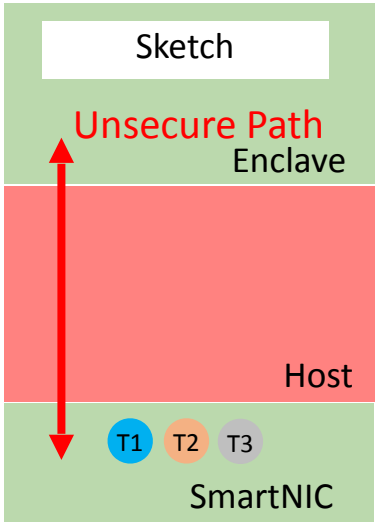

Packet Stream

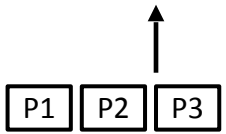

Challenge: Handling multi-threaded NICs

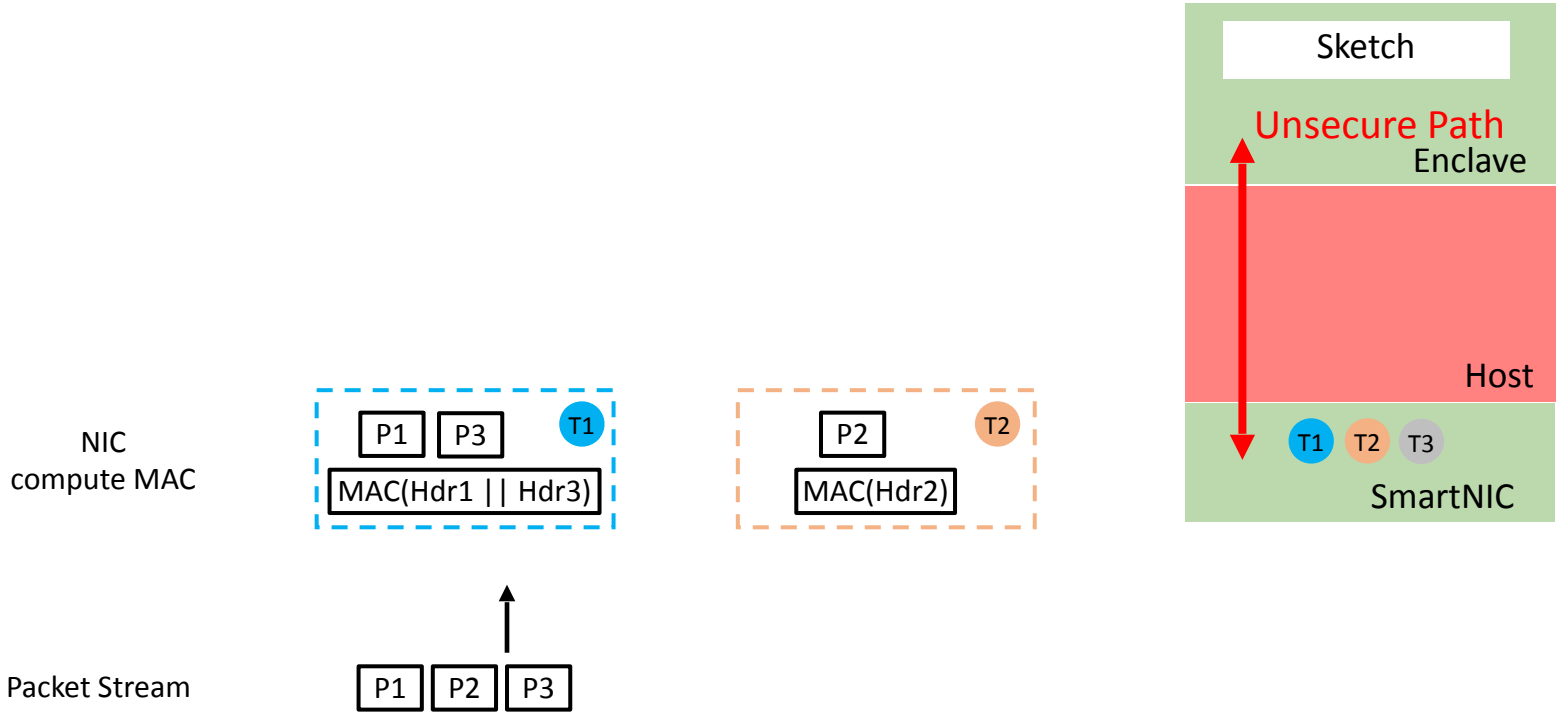

Challenge: Handling multi-threaded NICs

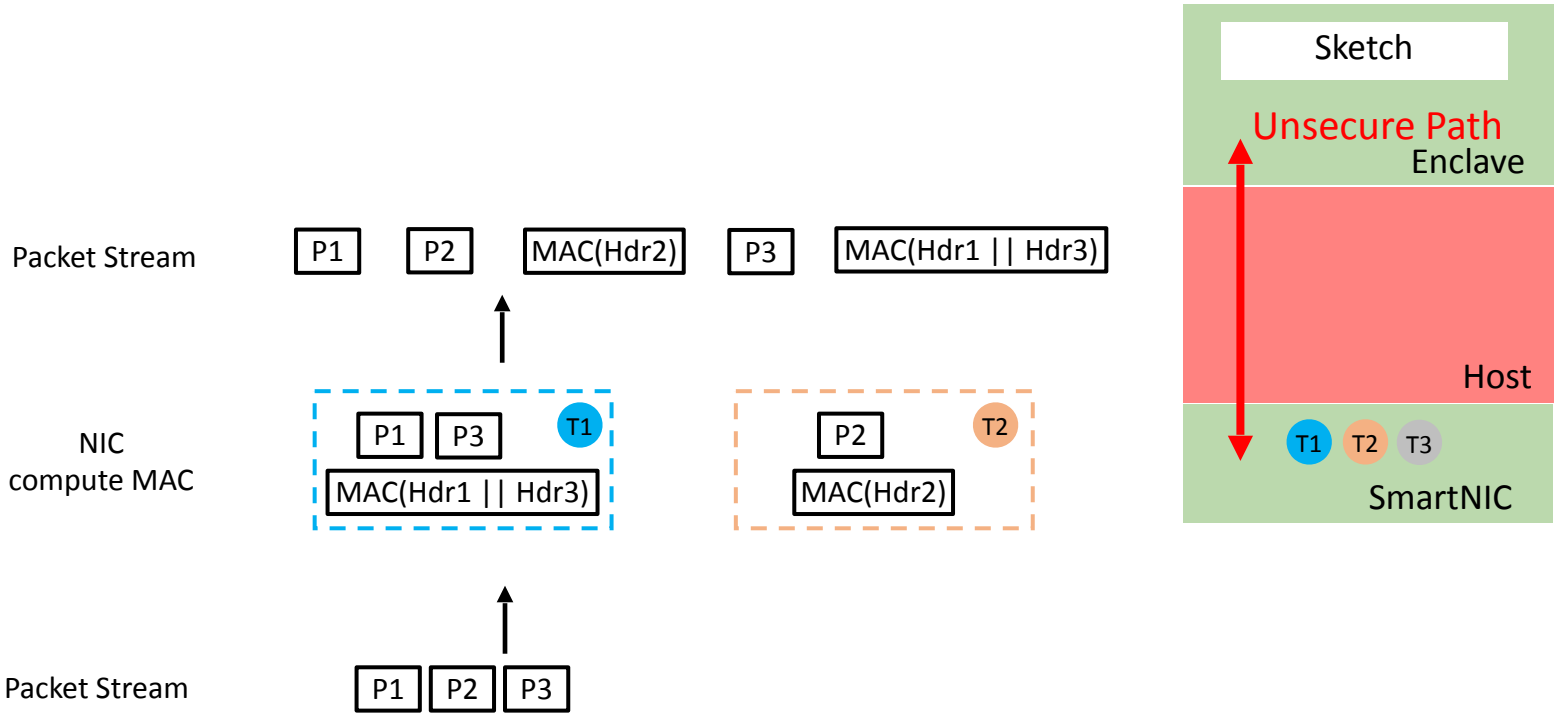

Challenge: Handling multi-threaded NICs

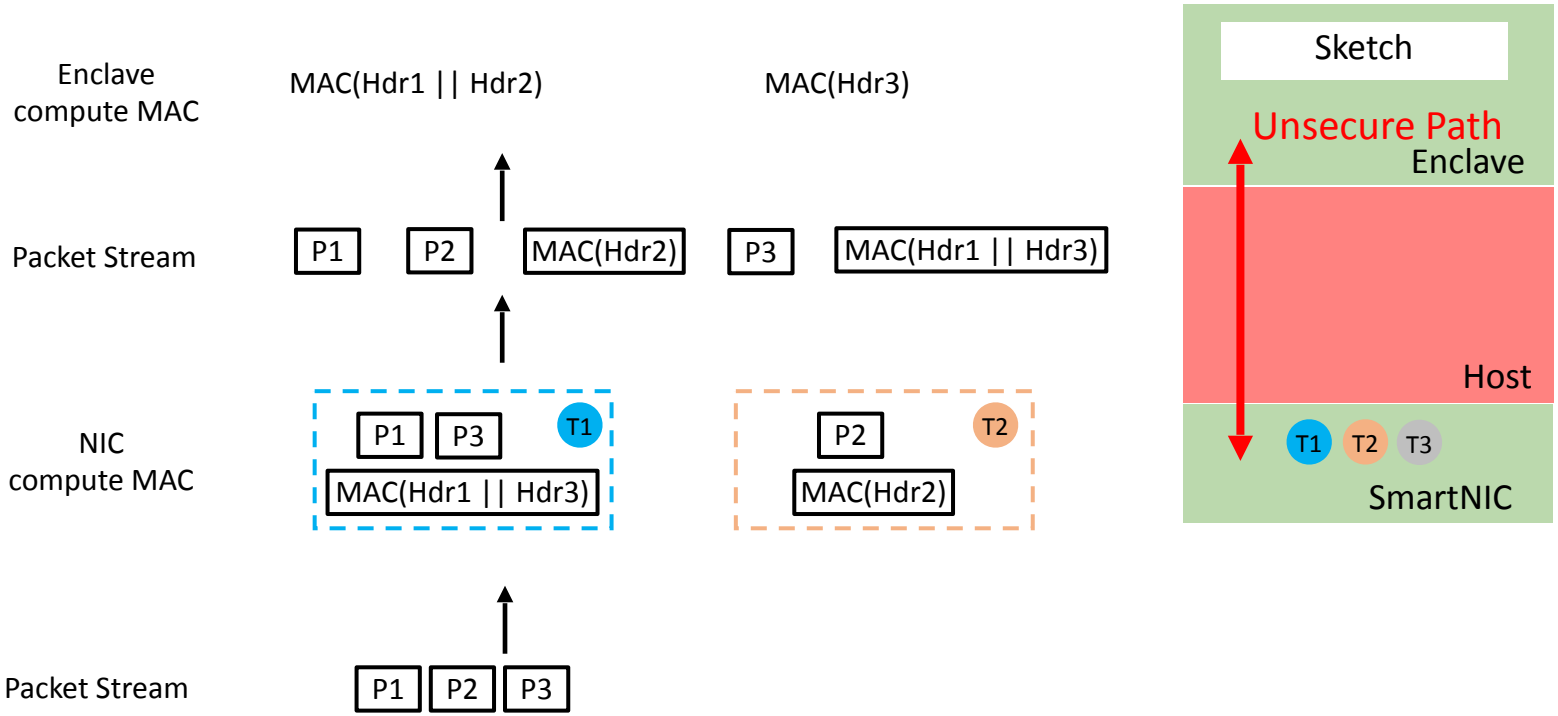

Challenge: Handling multi-threaded NICs

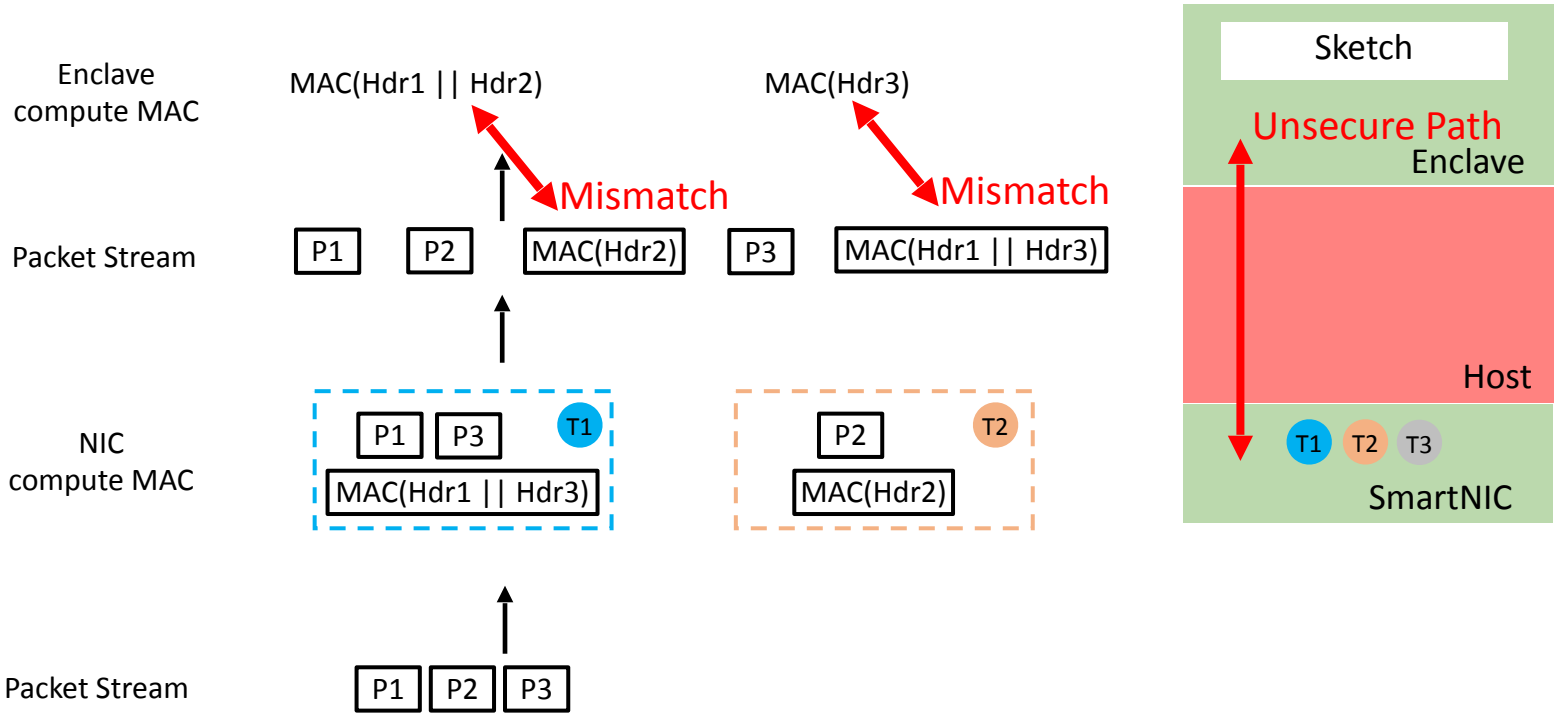

Solution: Tag packets to reconstruct substreams

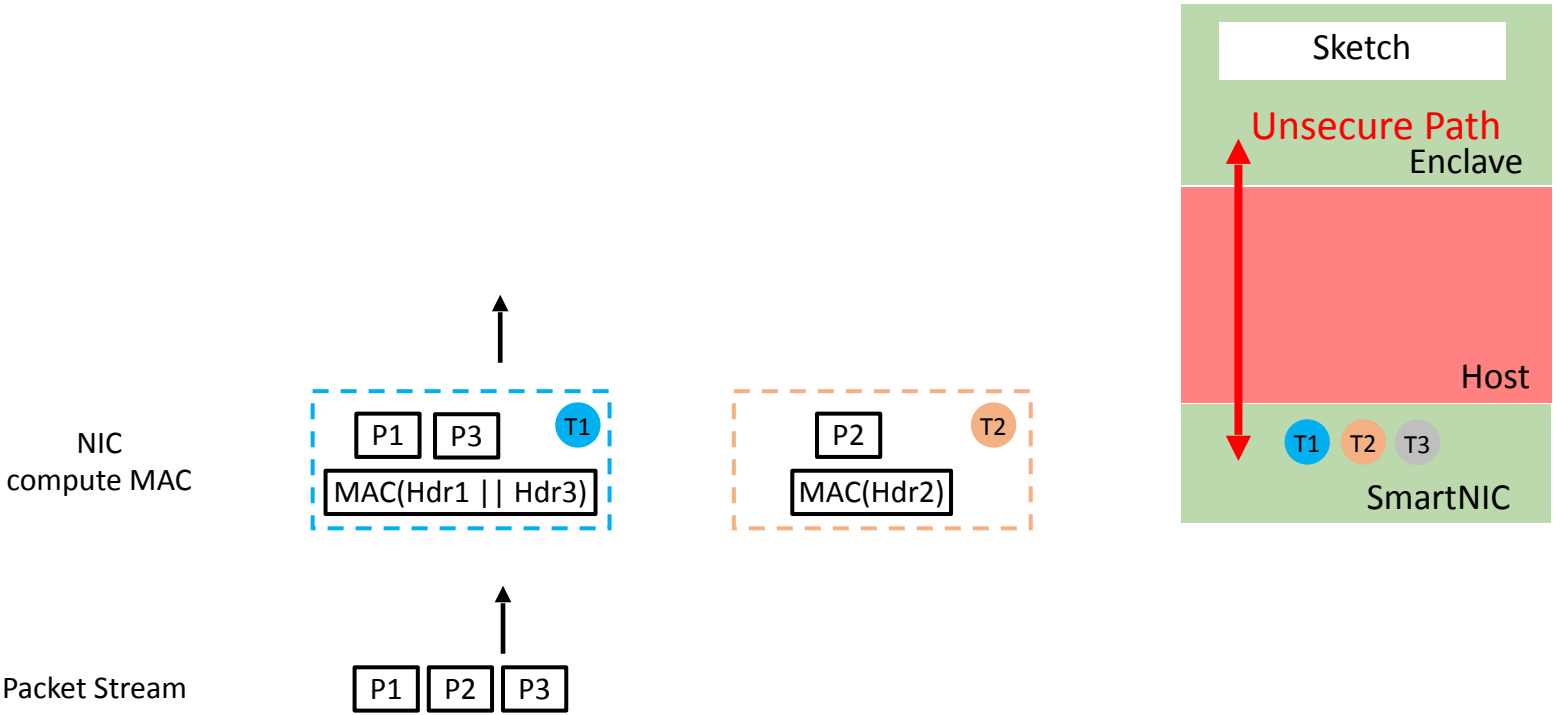

Solution: Tag packets to reconstruct substreams

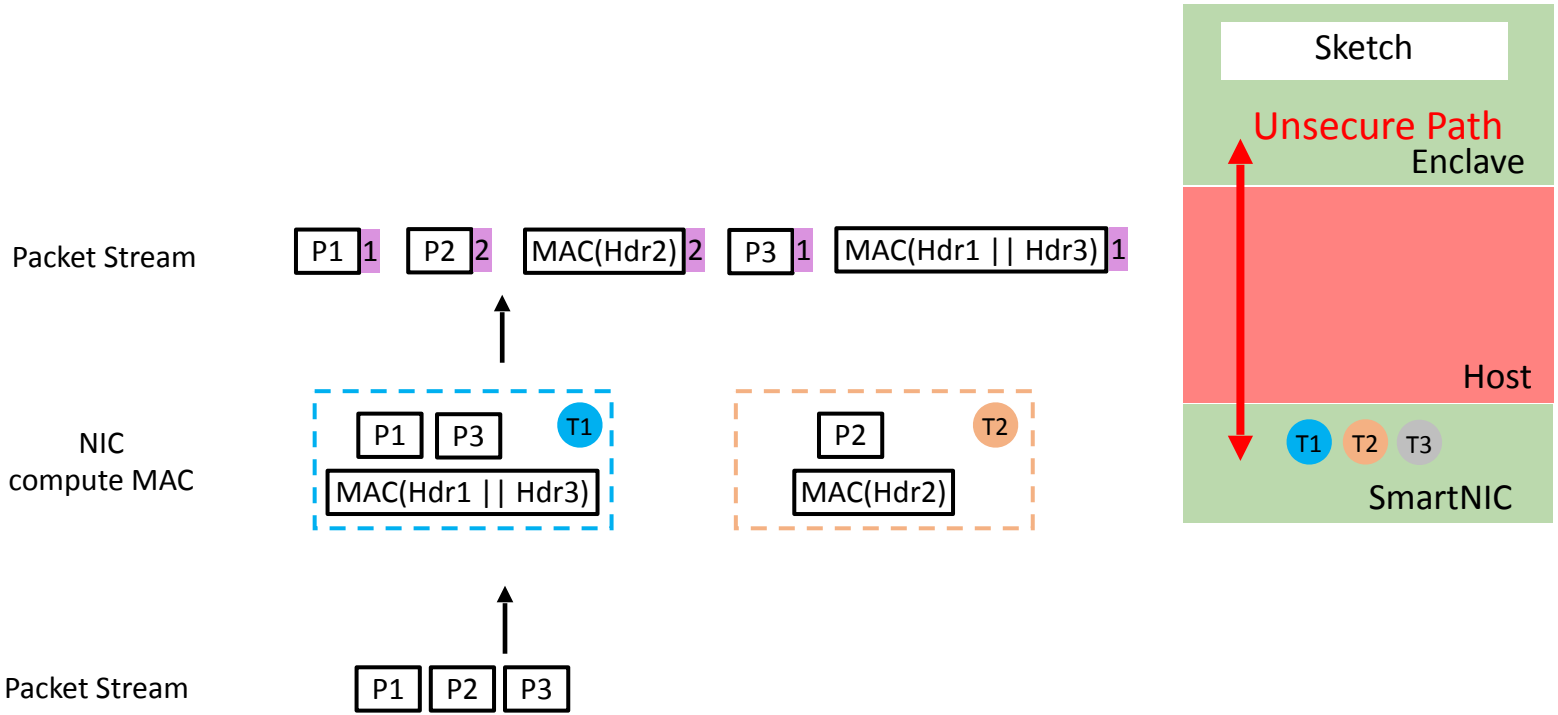

Solution: Tag packets to reconstruct substreams

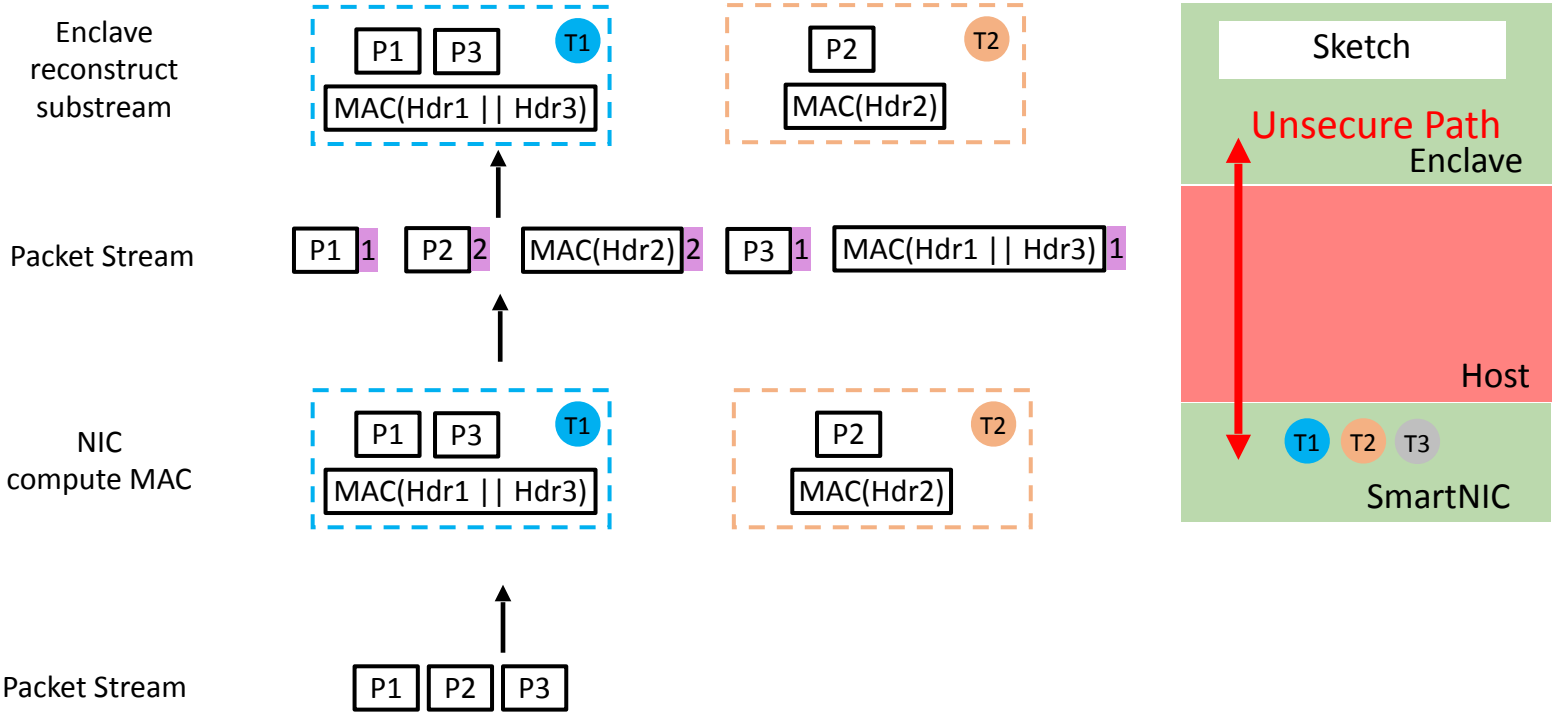

Talk Outline

- Motivation.
- Formulate requirements for trustworthy sketch-based telemetry.
- TrustSketch Design.
- Evaluation.

Evaluation: Security

Attack Type	Attack	Baseline (Unsafe)	Safebricks (Enclave Only)	Trustsketch
Compute	Modify runtime library			
Memory	Modify counter			
	Modify flow keys			
Input	Inject packets			
	Drop packets			
	Modify packet header			

Evaluation: Security

Attack Type	Attack	Baseline (Unsafe)	Safebricks (Enclave Only)	Trustsketch
Compute	Modify runtime library	✗	✓	✓
Memory	Modify counter	✗	✓	✓
	Modify flow keys	✗	✓	✓
Input	Inject packets	✗	✗	✓
	Drop packets	✗	✗	✓
	Modify packet header	✗	✗	✓

Evaluation: Performance

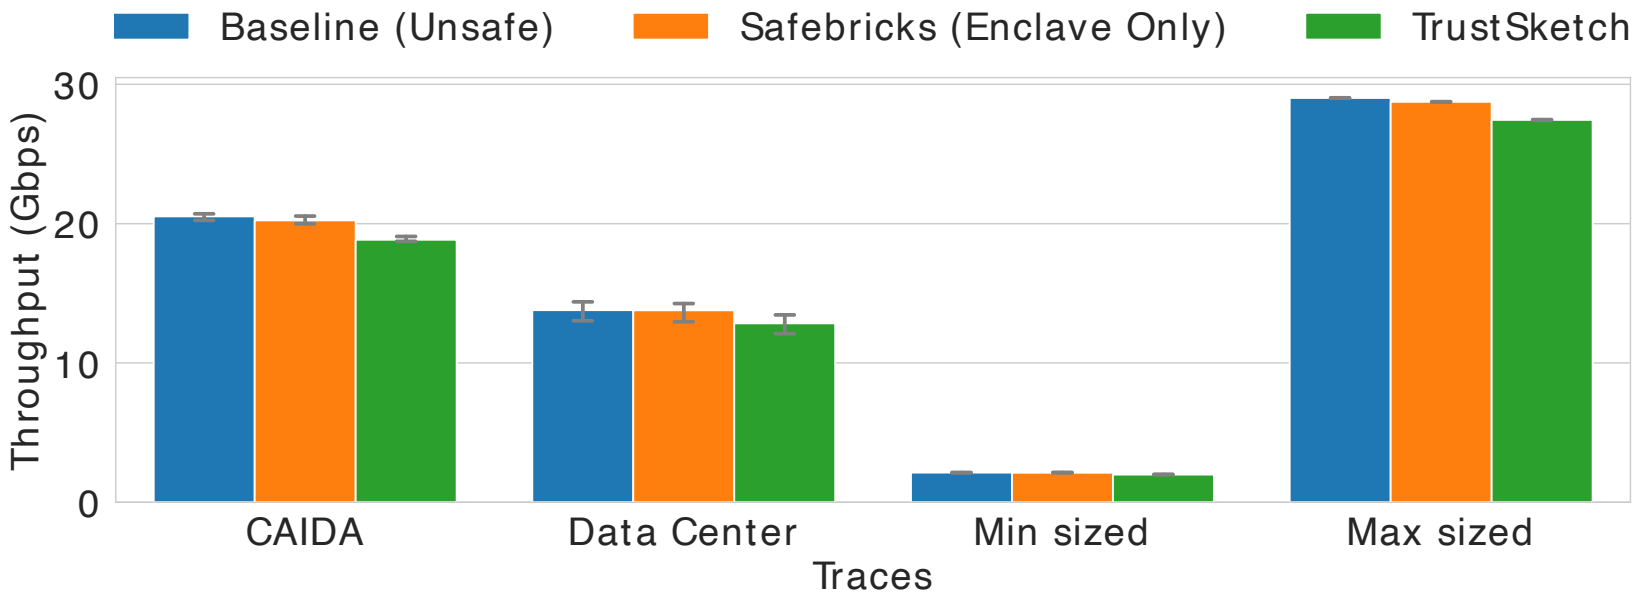

Compared to Baseline (unsafe), TrustSketch degrades throughput by 7%.

Evaluation: Accuracy

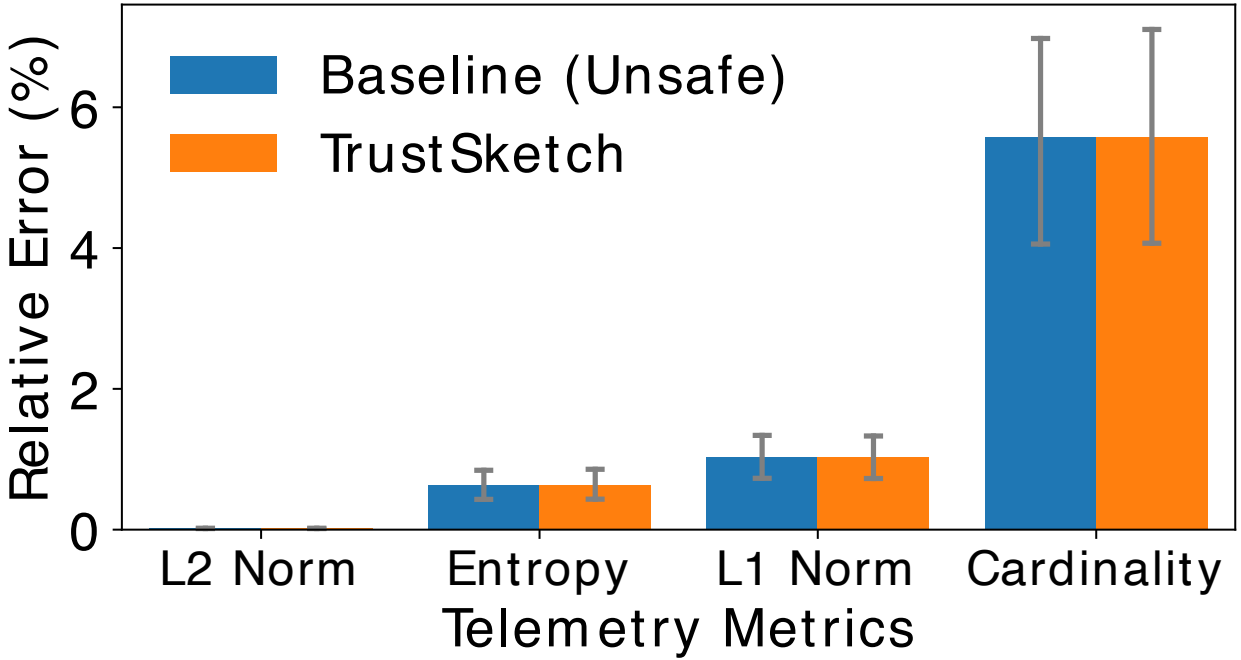

TrustSketch has the same accuracy as Baseline (Unsafe).

Summary

- Sketches are attractive for resource efficient telemetry in cloud.
- Existing architectures can be insecure if deployed naively.
- Contributions:
  - Formulate trustworthy sketch-based telemetry problem.
  - TrustSketch: based on enclave and SmartNIC.
  - Evaluation shows that TrustSketch is safe with low performance overhead.