A Double-Edged Sword: Security Threats and Opportunities in One-Sided Network Communication

Shin-Yeh Tsai
Yiying Zhang

WukLab

PURDUE UNIVERSITY

Traditional (Two-sided communication)

User A

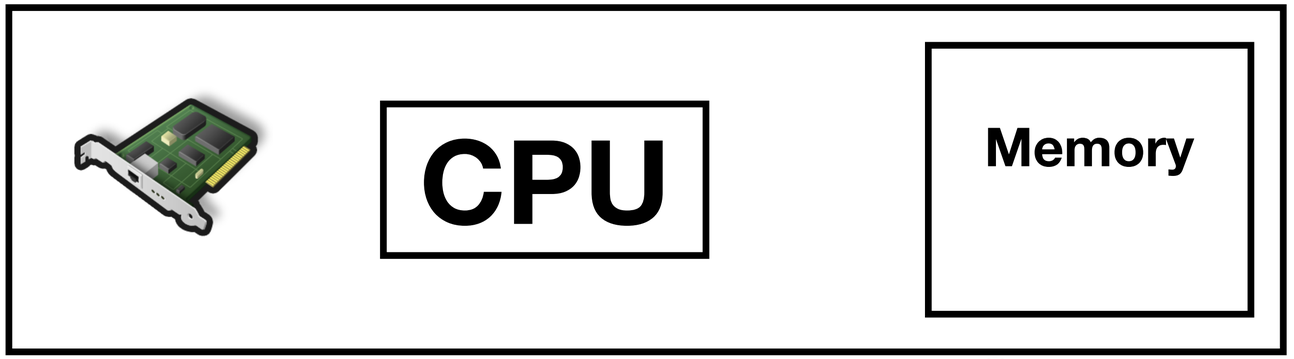

Traditional (Two-sided communication)

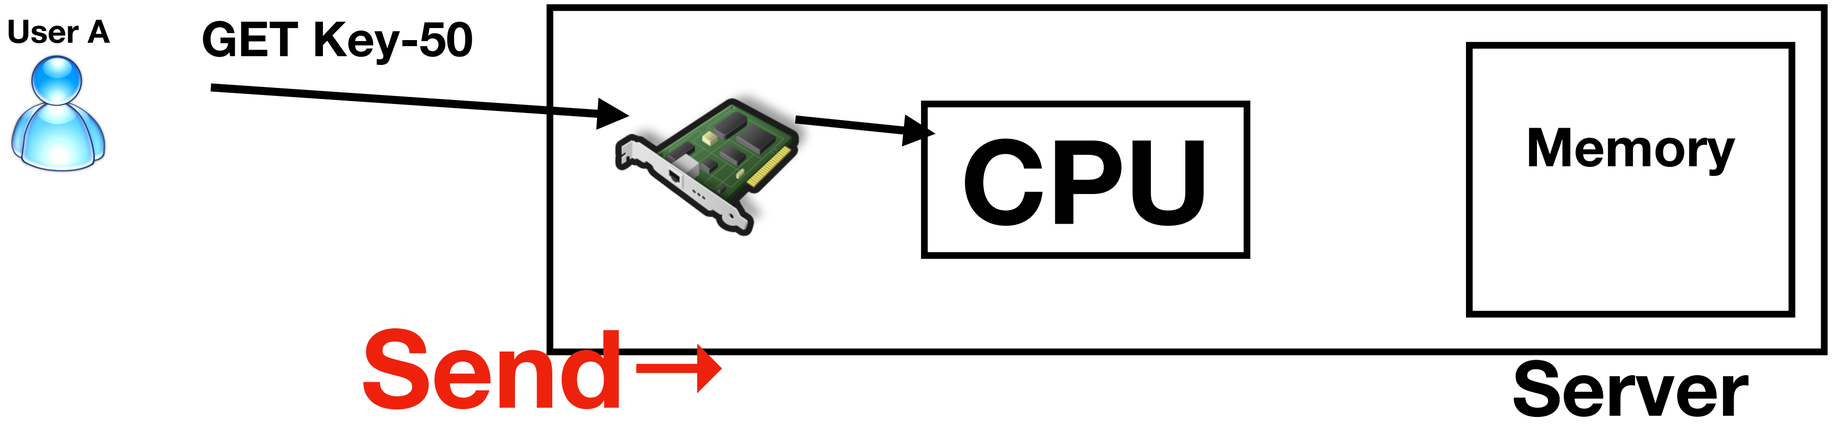

Traditional (Two-sided communication)

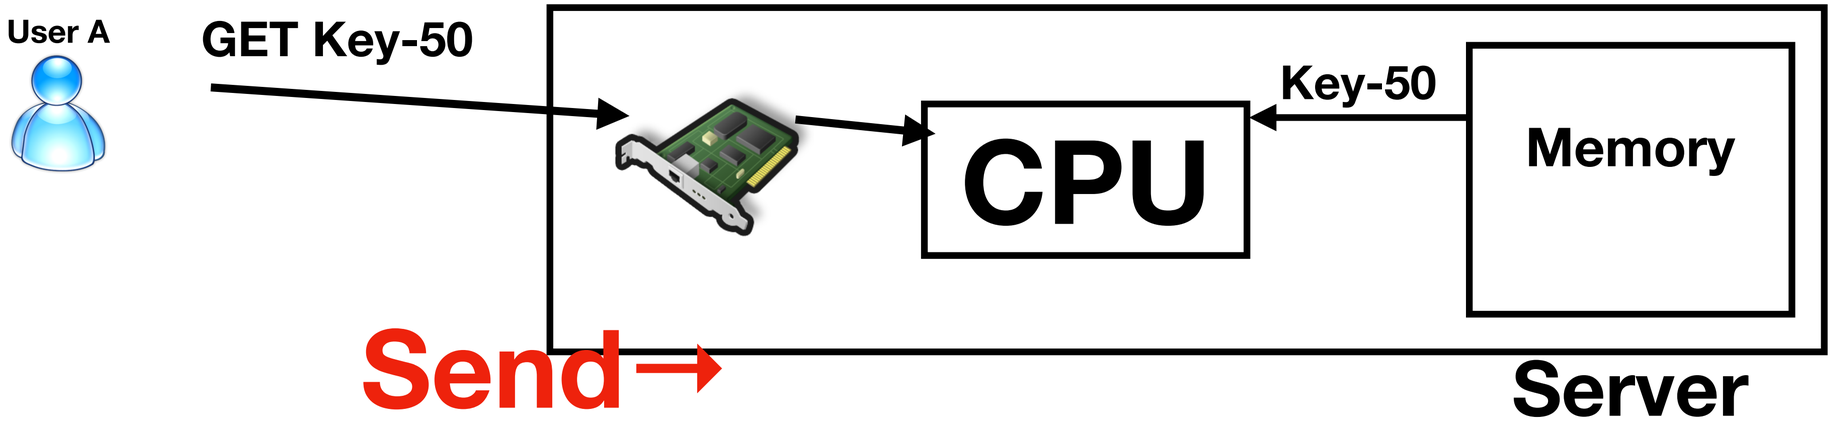

Traditional (Two-sided communication)

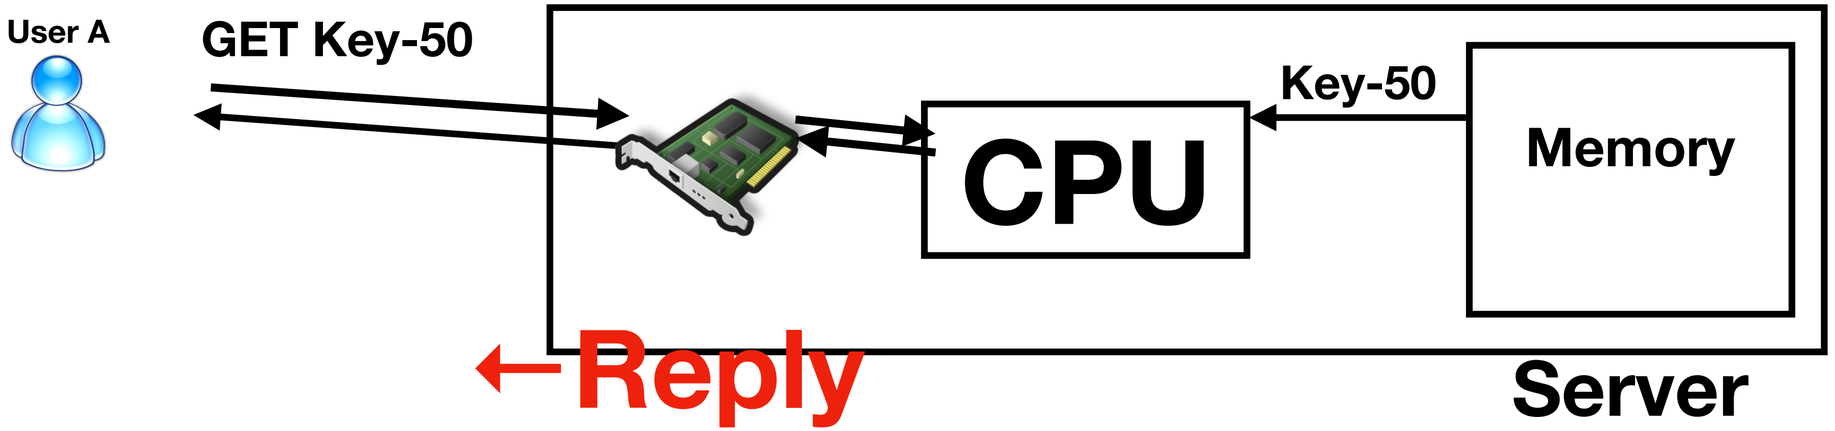

Traditional (Two-sided communication)

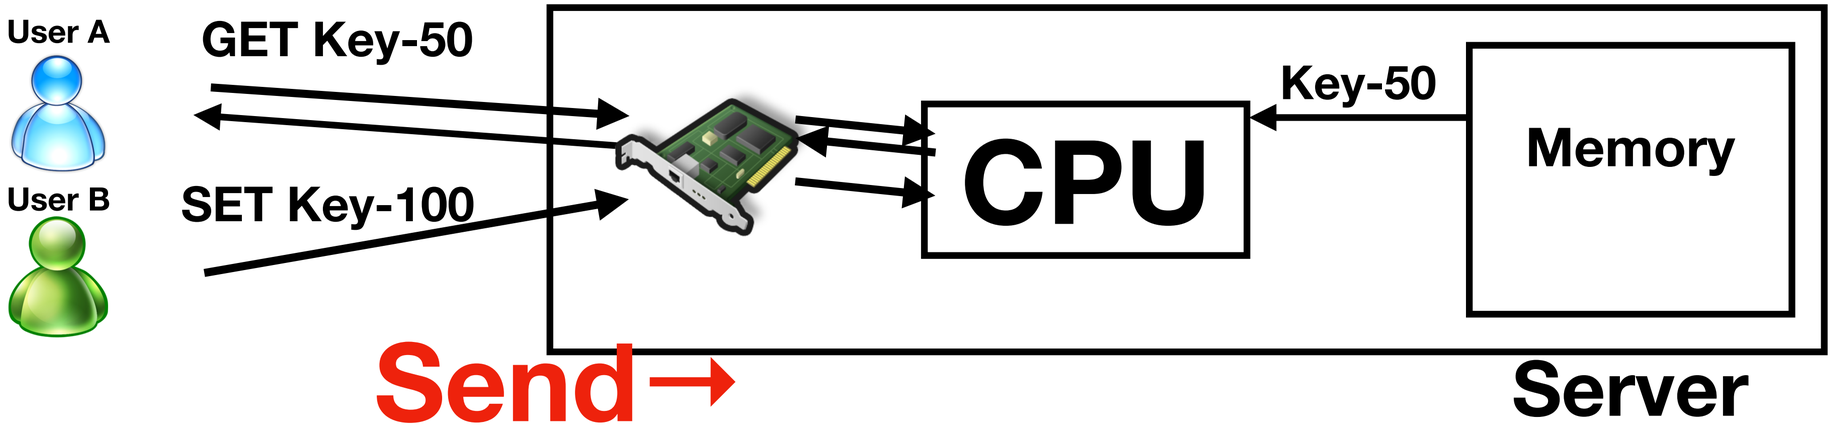

Traditional (Two-sided communication)

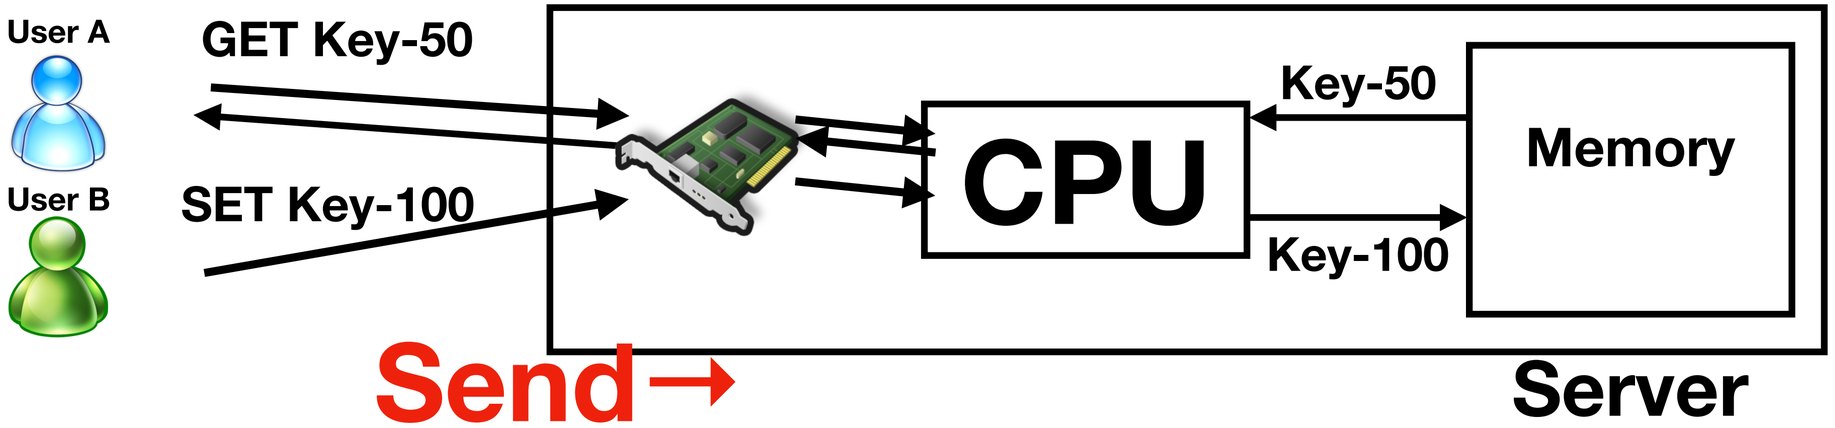

Send→

Traditional (Two-sided communication)

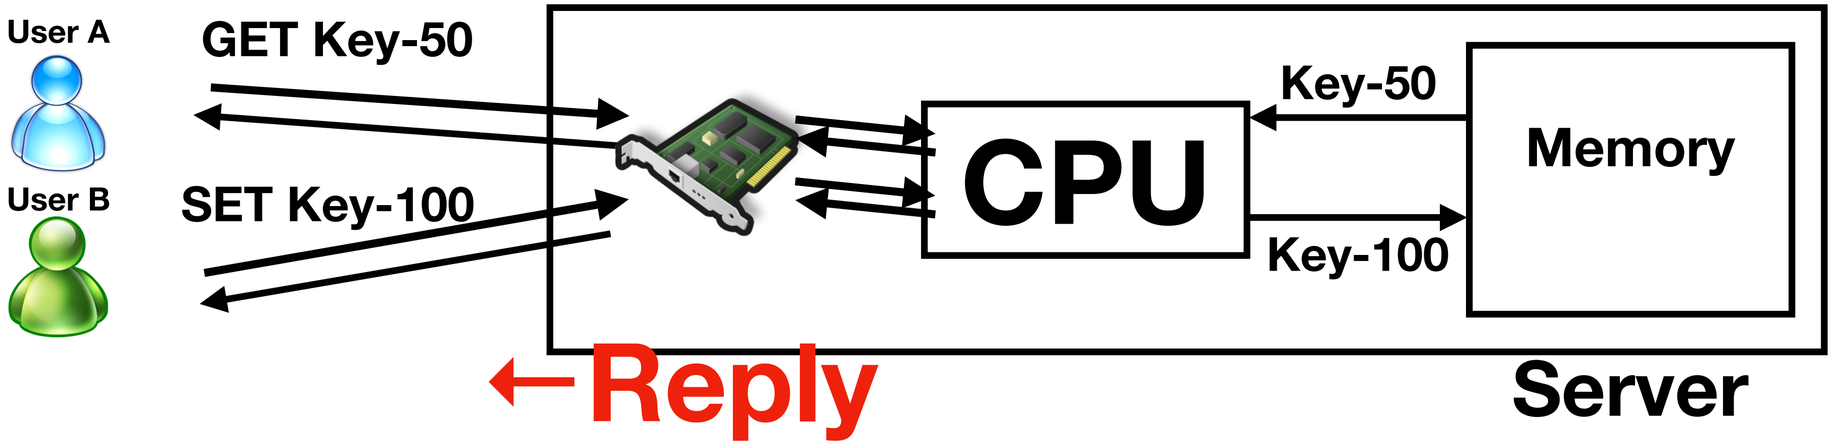

Traditional (Two-sided communication)

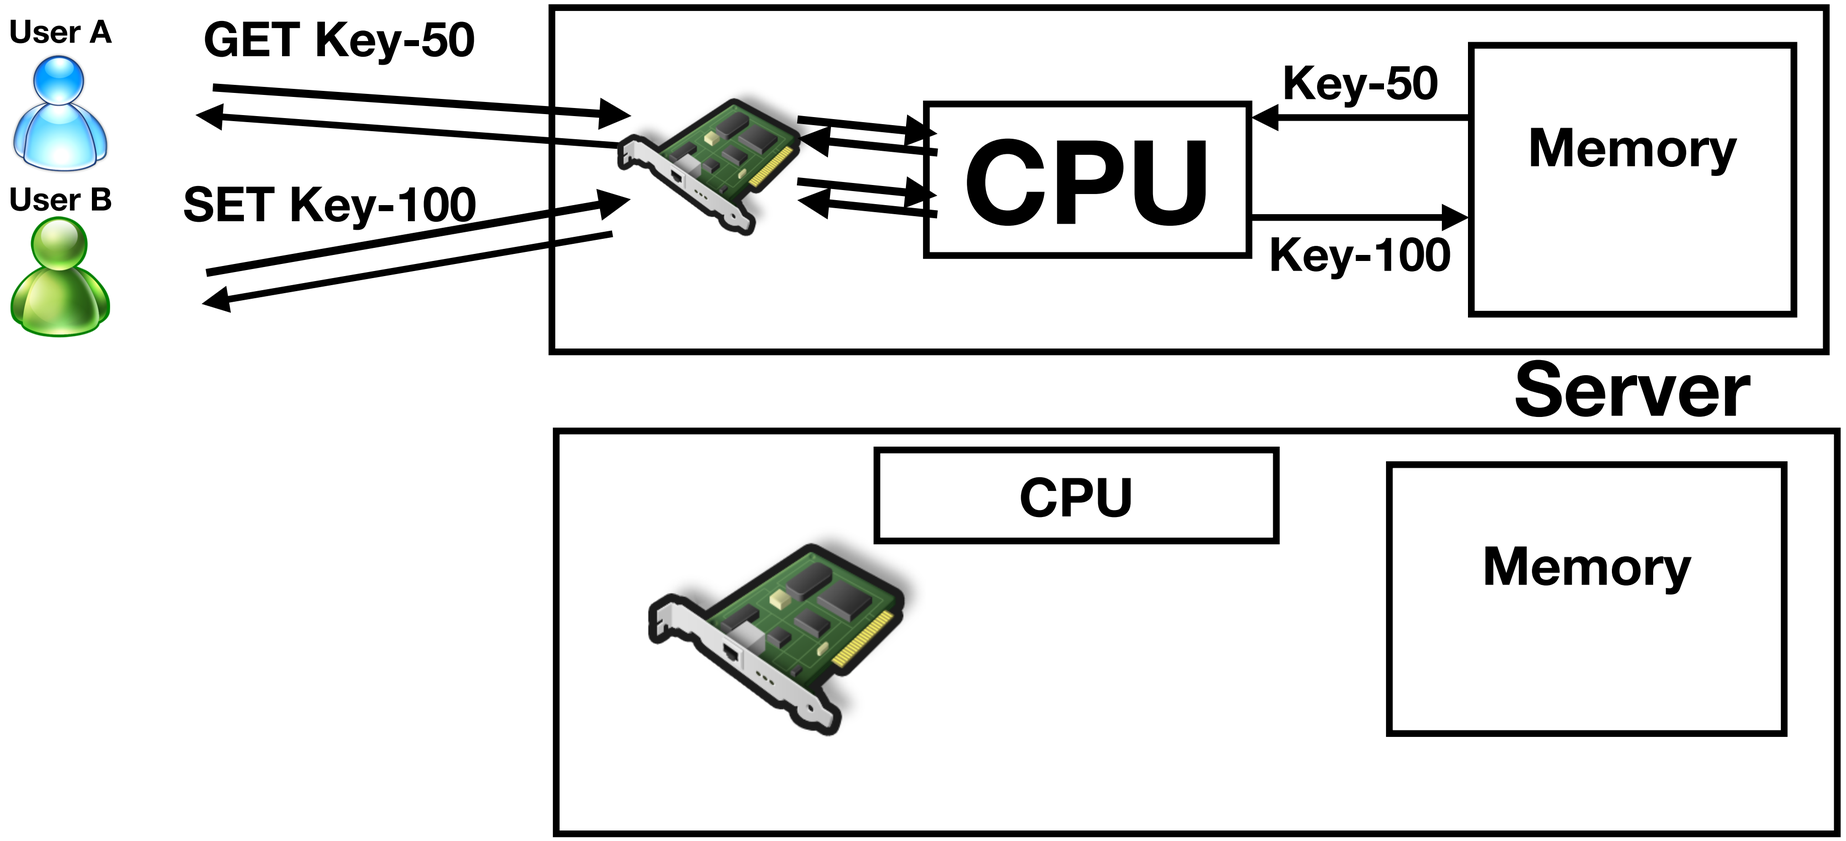

One-sided communication

Traditional (Two-sided communication)

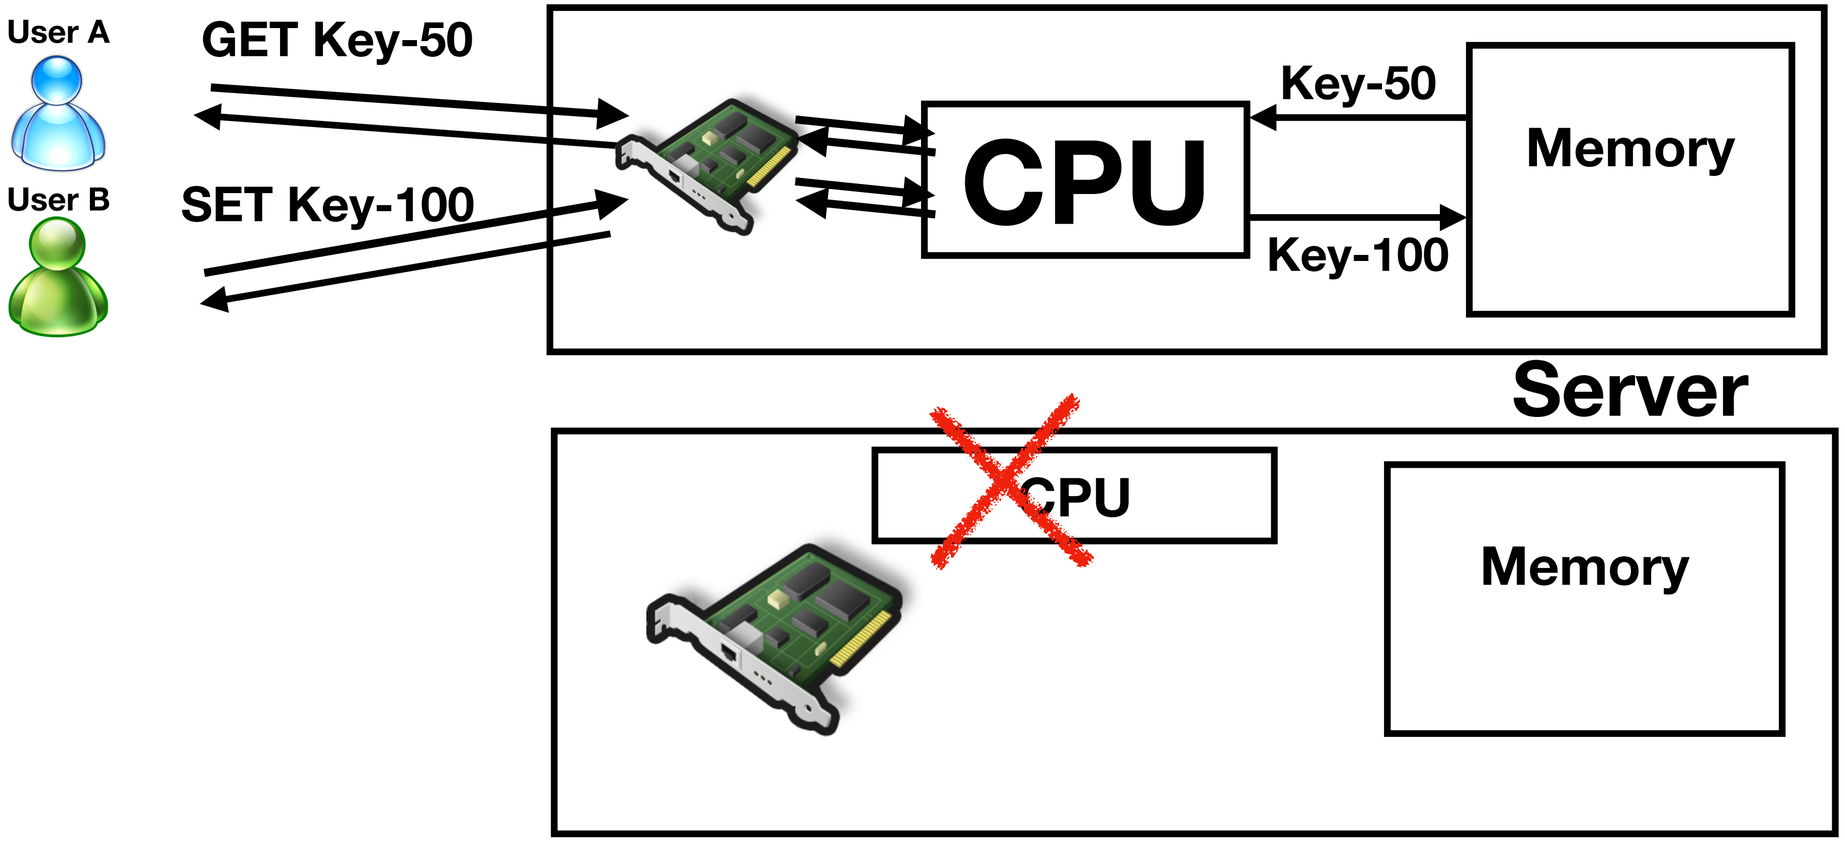

One-sided communication

Traditional (Two-sided communication)

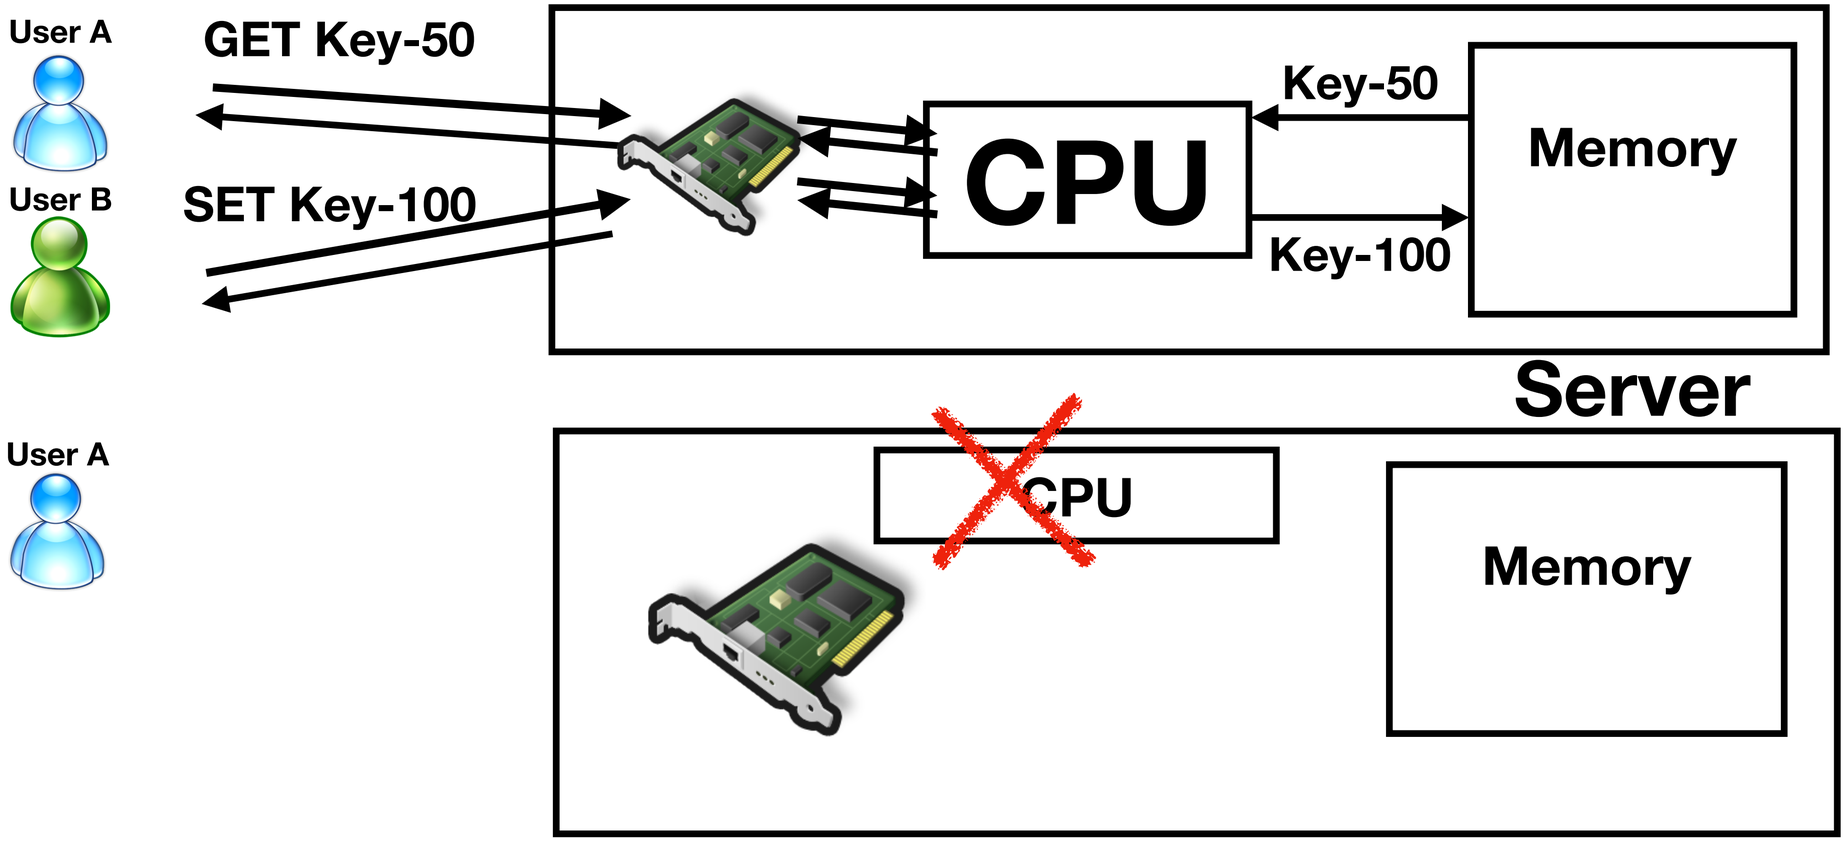

One-sided communication

Traditional (Two-sided communication)

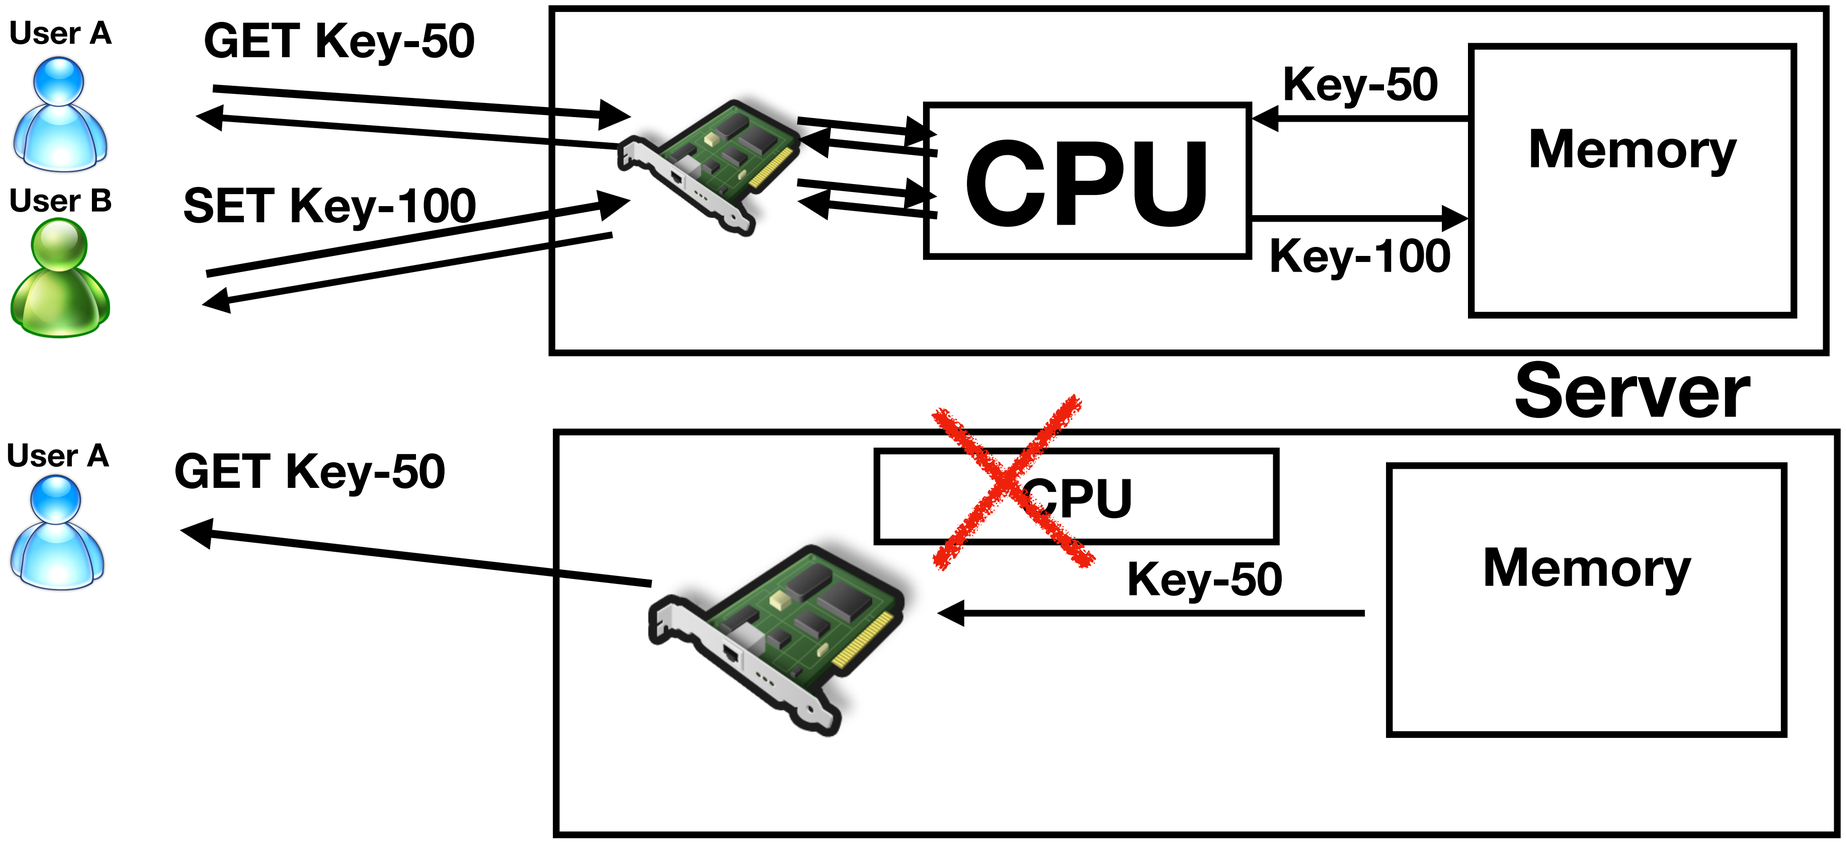

One-sided communication

Traditional (Two-sided communication)

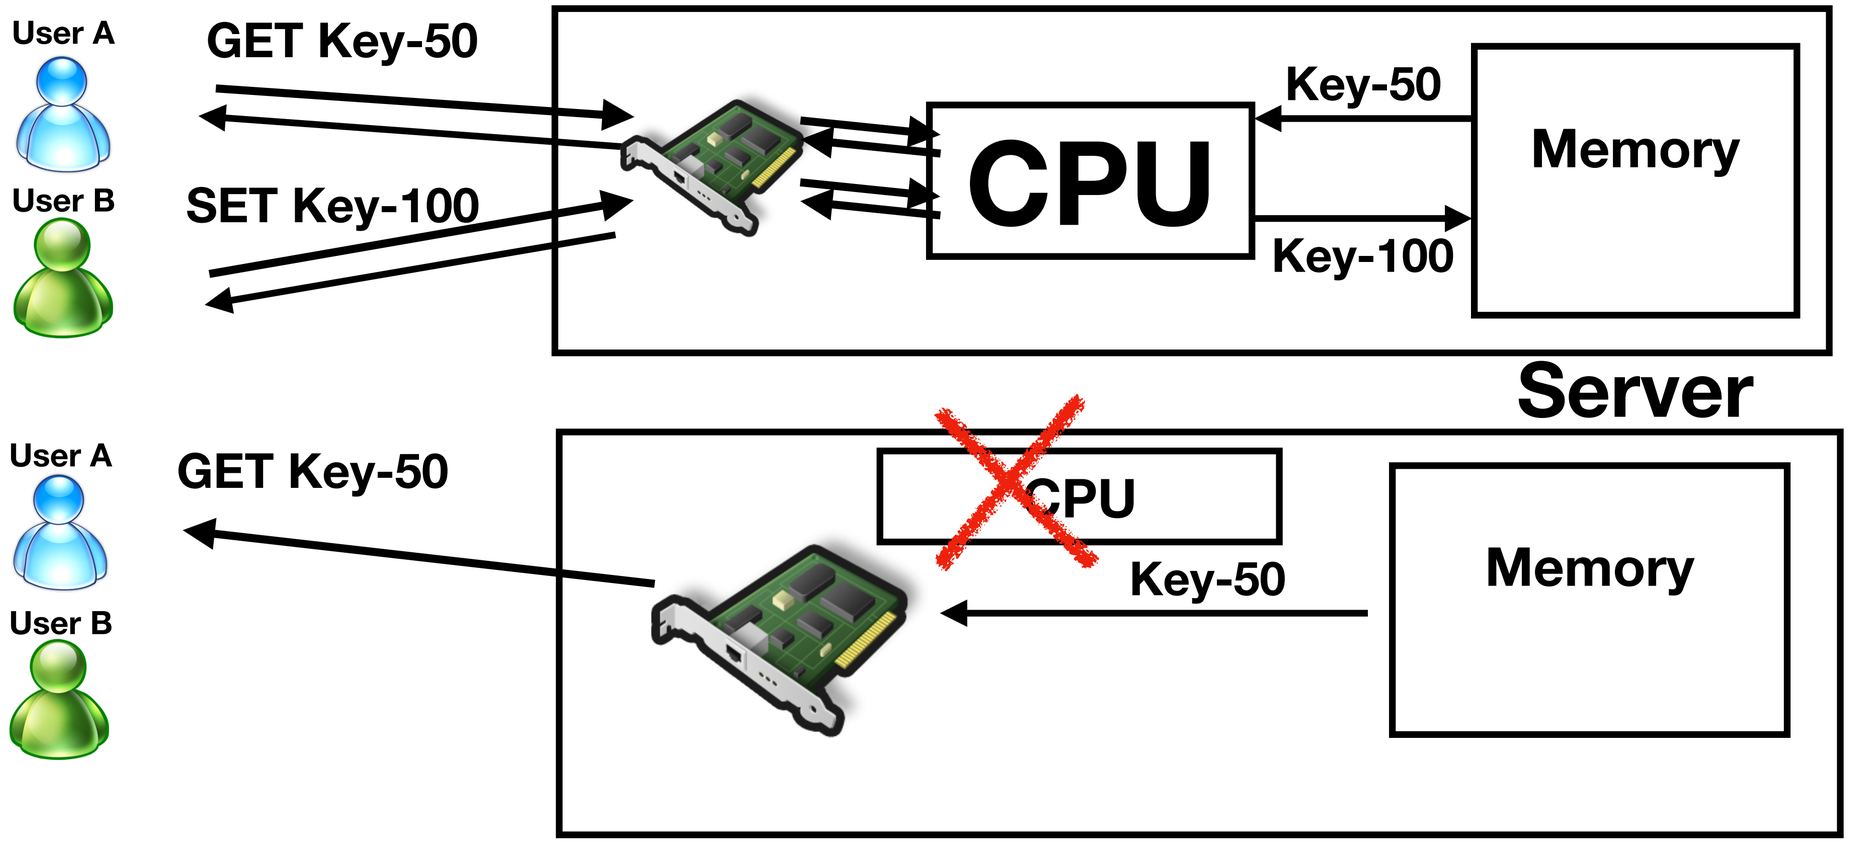

One-sided communication

Traditional (Two-sided communication)

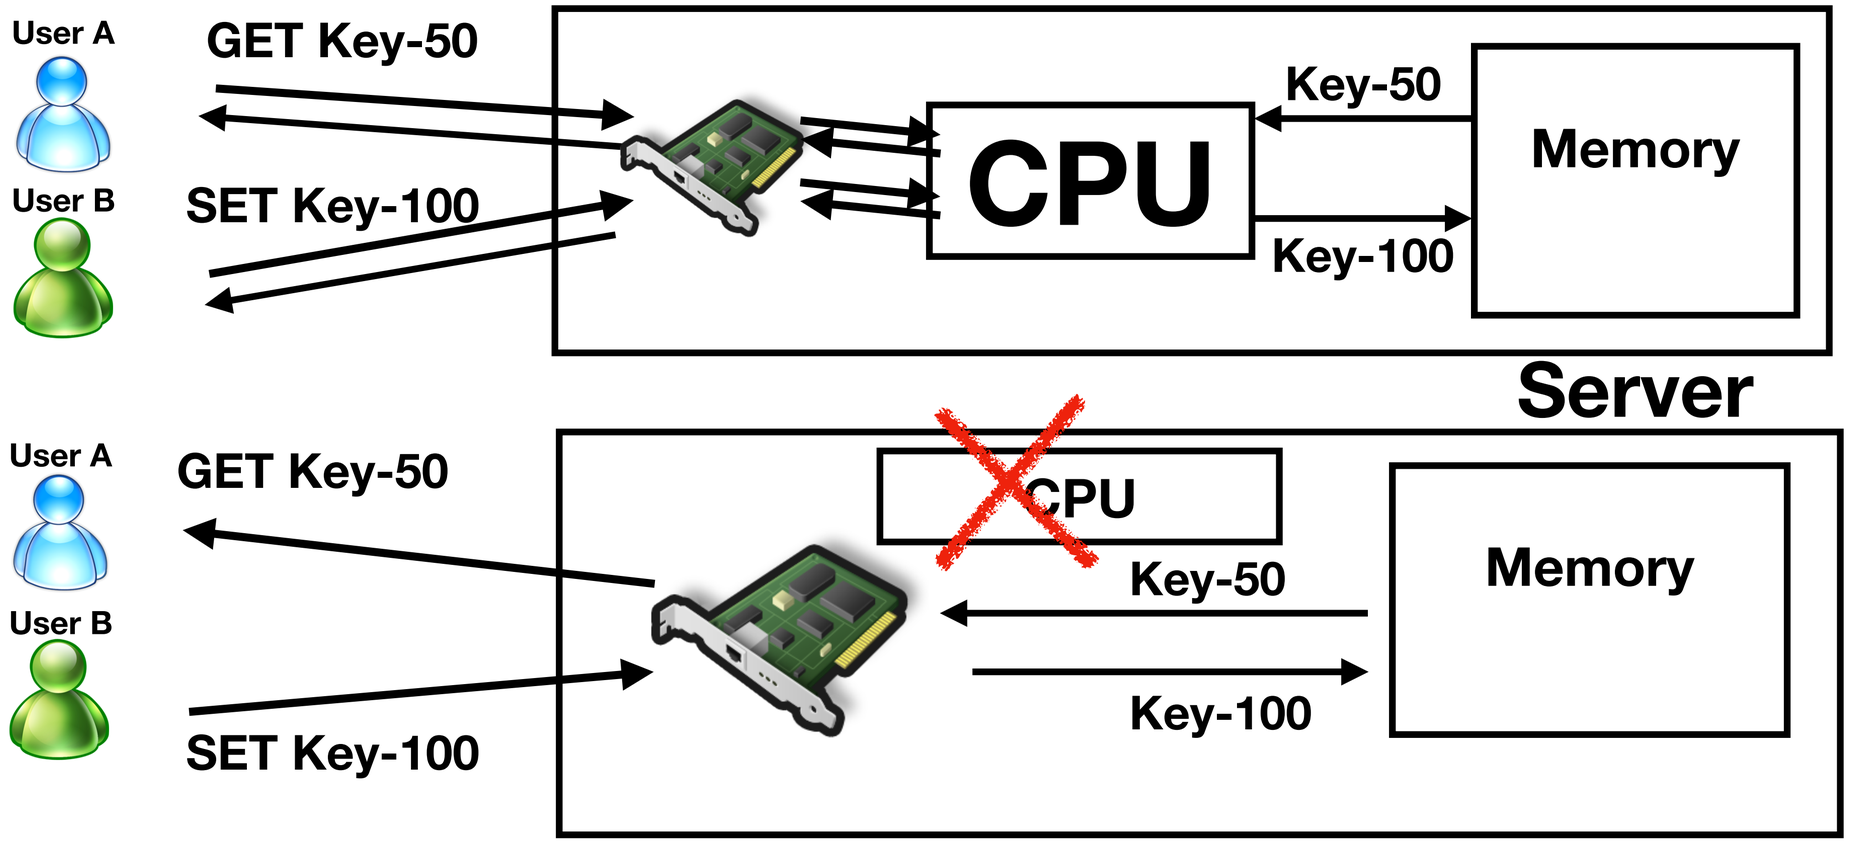

One-sided communication

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

Pilaf
[ATC ’13]

APUS
[SoCC ’17]

Hotpot
[SoCC ’17]

FaRM
[NSDI ’14]

Octopus
[ATC ’17]

Orion
[FAST ’19]

Wukong
[OSDI ’16]

Storm
[SYSTOR '19]

HERD
[SIGCOMM ’14]

Cell
[ATC ’16]

NAM-DB
[VLDB ’17]

DRTM+H
[OSDI ’18]

DrTM
[SOSP ’15]

KV-Direct
[SOSP ’17]

FaSST
[OSDI ’16]

LITE
[SOSP ’17]

FaRM + Xact
[SOSP ’15]

RSI
[VLDB ’16]

Mojim
[ASPLOS ’15]

DrTM+R
[EuroSys ’16]

Performance

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

Hotpot
[SoCC ’17]

Pilaf
[ATC ’13]

APUS
[SoCC ’17]

FaRM
[NSDI ’14]

Octopus
[ATC ’17]

Orion
[FAST ’19]

Wukong
[OSDI ’16]

Storm
[SYSTOR '19]

HERD
[SIGCOMM ’14]

Cell
[ATC ’16]

NAM-DB
[VLDB ’17]

DRTM+H
[OSDI ’18]

DrTM
[SOSP ’15]

KV-Direct
[SOSP ’17]

FaSST
[OSDI ’16]

LITE
[SOSP ’17]

FaRM + Xact
[SOSP ’15]

RSI
[VLDB ’16]

Mojim
[ASPLOS ’15]

DrTM+R
[EuroSys ’16]

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

Performance

Scalability

Hotpot
[SoCC ’17]

Wukong
[OSDI ’16]

HERD
[SIGCOMM ’14]

NAM-DB
[VLDB ’17]

DrTM
[SOSP ’15]

KV-Direct
[SOSP ’17]

FaSST
[OSDI ’16]

LITE
[SOSP ’17]

FaRM + Xact
[SOSP ’15]

RSI
[VLDB ’16]

Mojim
[ASPLOS ’15]

DrTM+R
[EuroSys ’16]

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

Performance

Scalability

Usability

Hotpot
[SoCC ’17]

Wukong
[OSDI ’16]

HERD
[SIGCOMM ’14]

NAM-DB
[VLDB ’17]

DrTM
[SOSP ’15]

FaSST
[OSDI ’16]

FaRM + Xact
[SOSP ’15]

RSI
[VLDB ’16]

Mojim
[ASPLOS ’15]

DrTM+R
[EuroSys ’16]

RDMA

Omni-Path

Gen-Z

NVMeOF

GPUDirect

Performance

Scalability

Usability

What about Security?

Hotpot
[SoCC '17]

Wukong
[OSDI '16]

HERD
[SIGCOMM '14]

NAM-DB
[VLDB '17]

DrTM
[SOSP '15]

FaRM + Xact
[SOSP '15]

FaSST
[OSDI '16]

RSI
[VLDB '16]

Outline

- Introduction and Background
- Vulnerabilities in One-Sided Communication
- Vulnerabilities in One-Sided Hardware
- Opportunities in One-Sided Communication
- Conclusion

Vulnerability 1: Lack of Accountability

- **WRITE accountability**

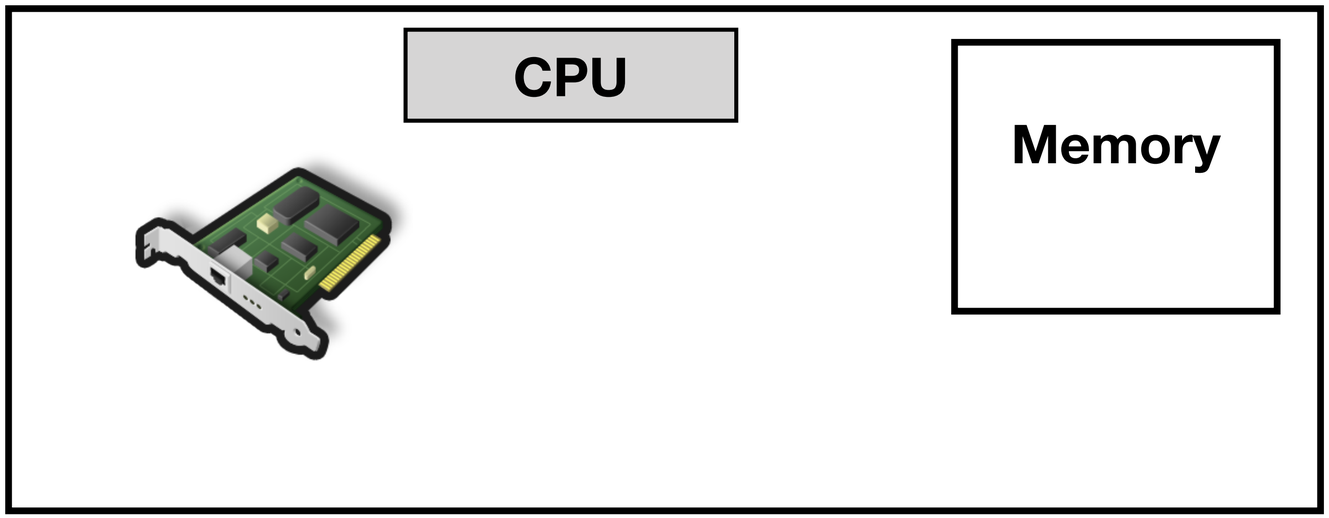

Vulnerability 1: Lack of Accountability

- **WRITE accountability**

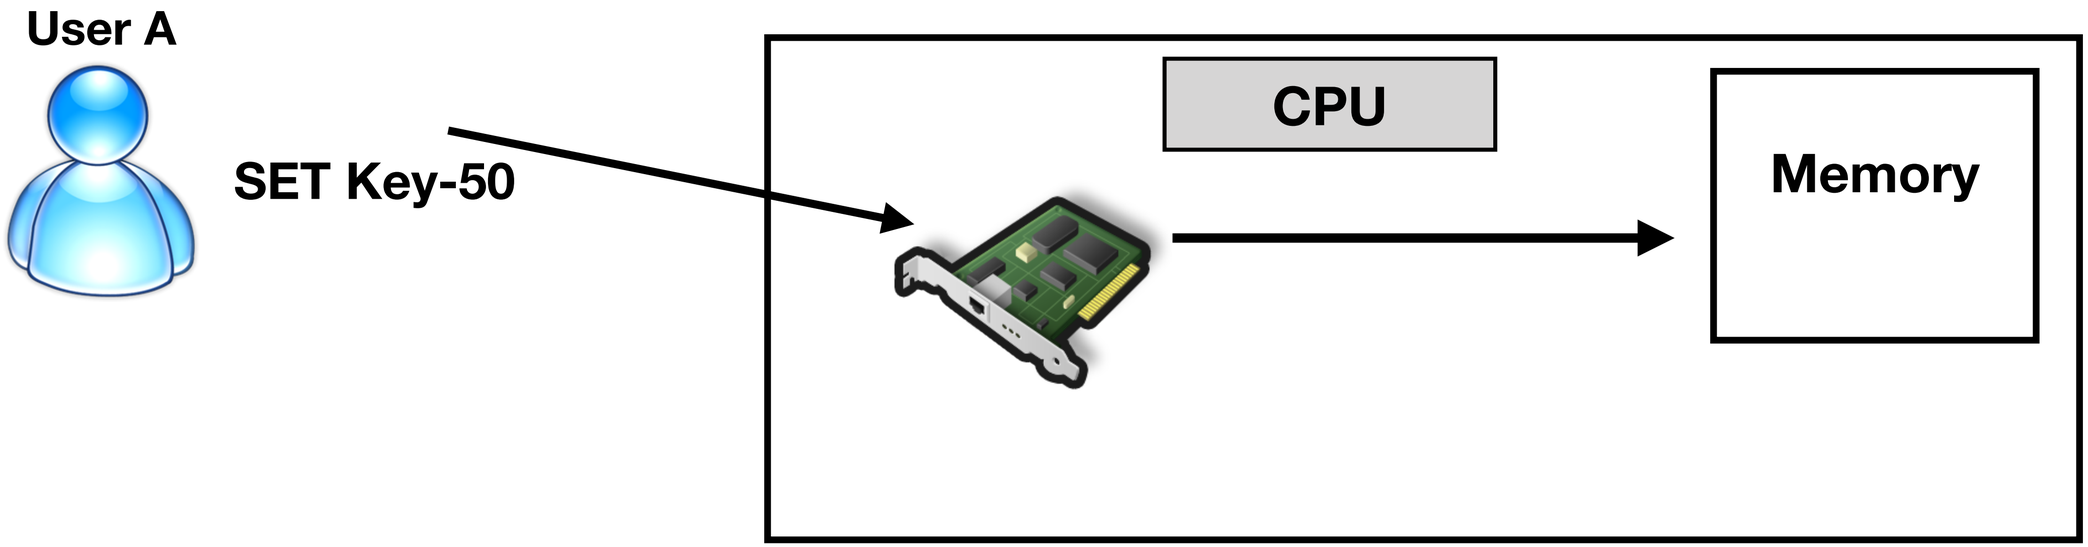

Vulnerability 1: Lack of Accountability

- **WRITE accountability**

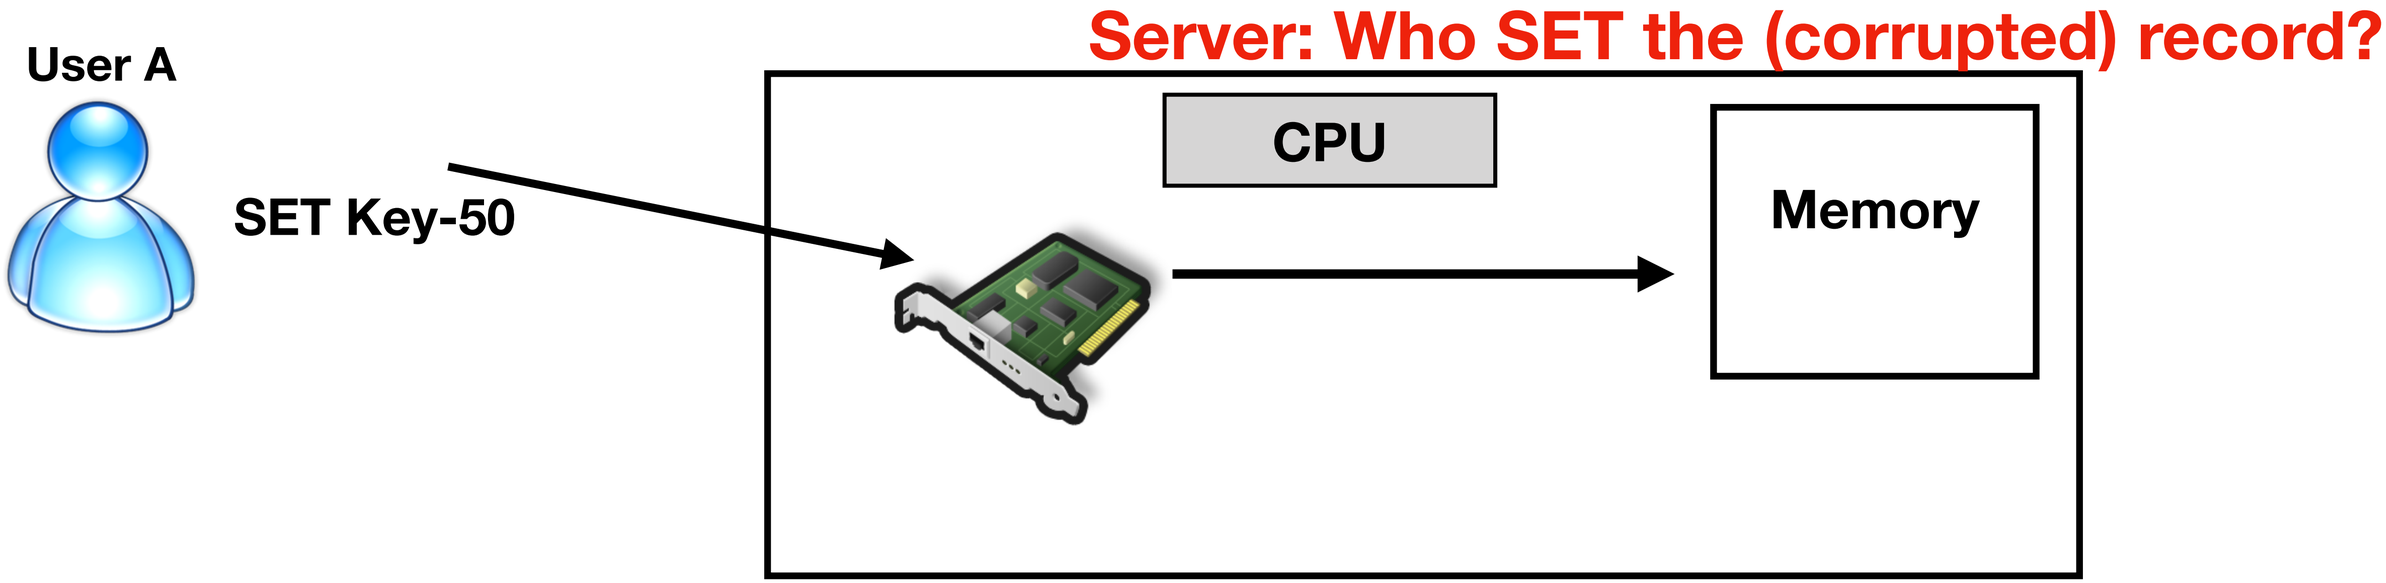

Vulnerability 1: Lack of Accountability

- **WRITE accountability**
- **READ accountability**

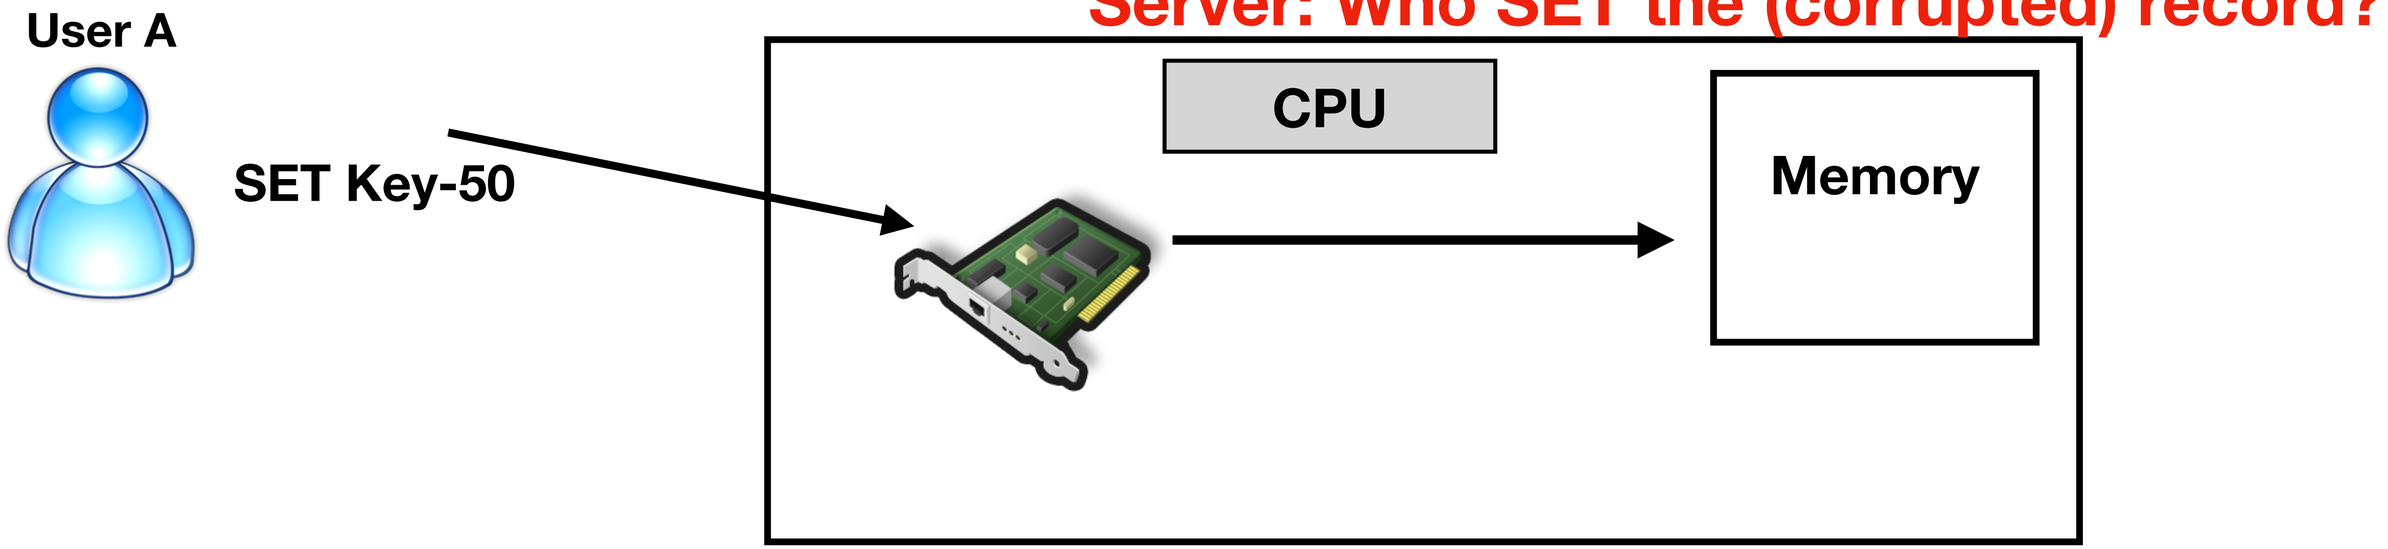

Vulnerability 1: Lack of Accountability

- **WRITE accountability**
- **READ accountability**

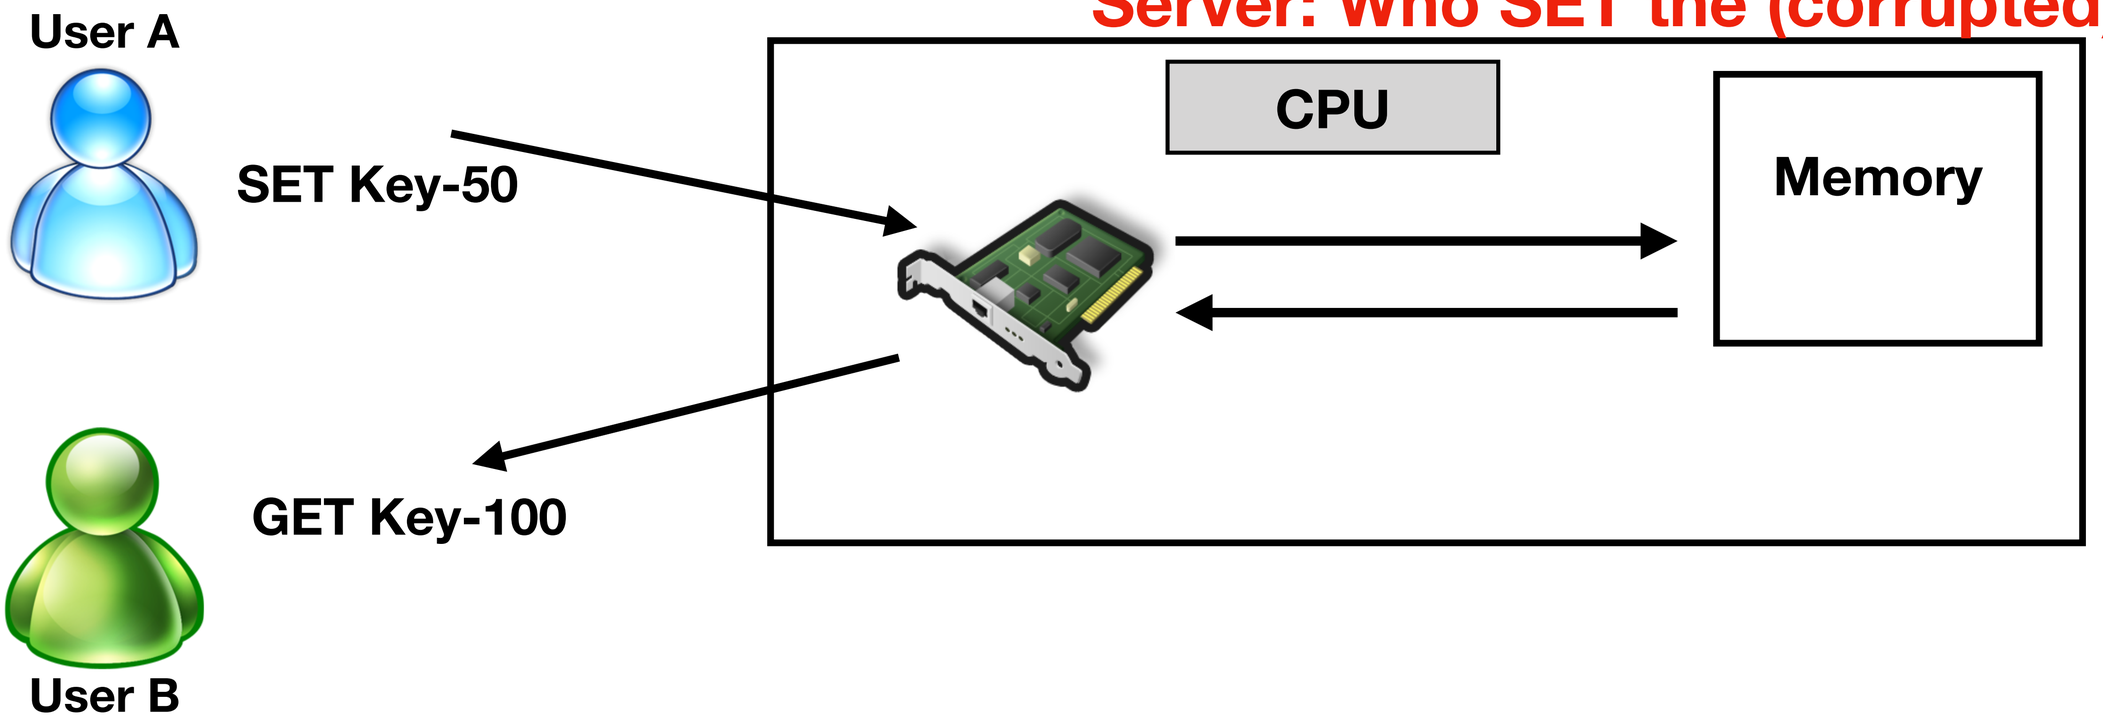

Vulnerability 1: Lack of Accountability

- **WRITE accountability**
- **READ accountability**

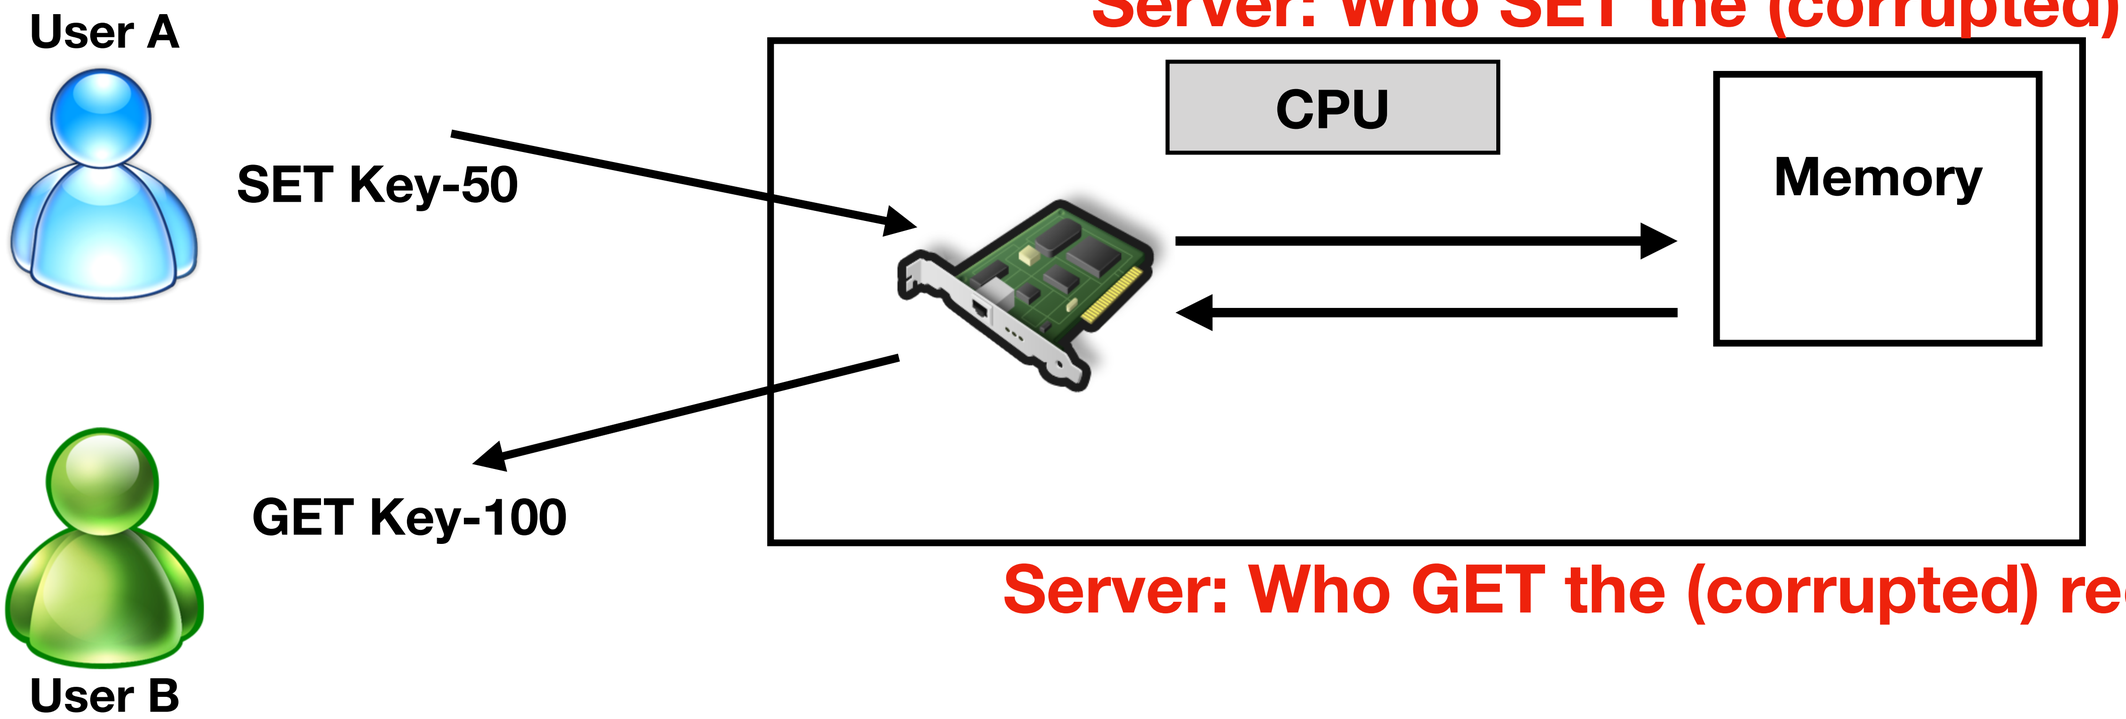

Vulnerability 2: Denial of Service

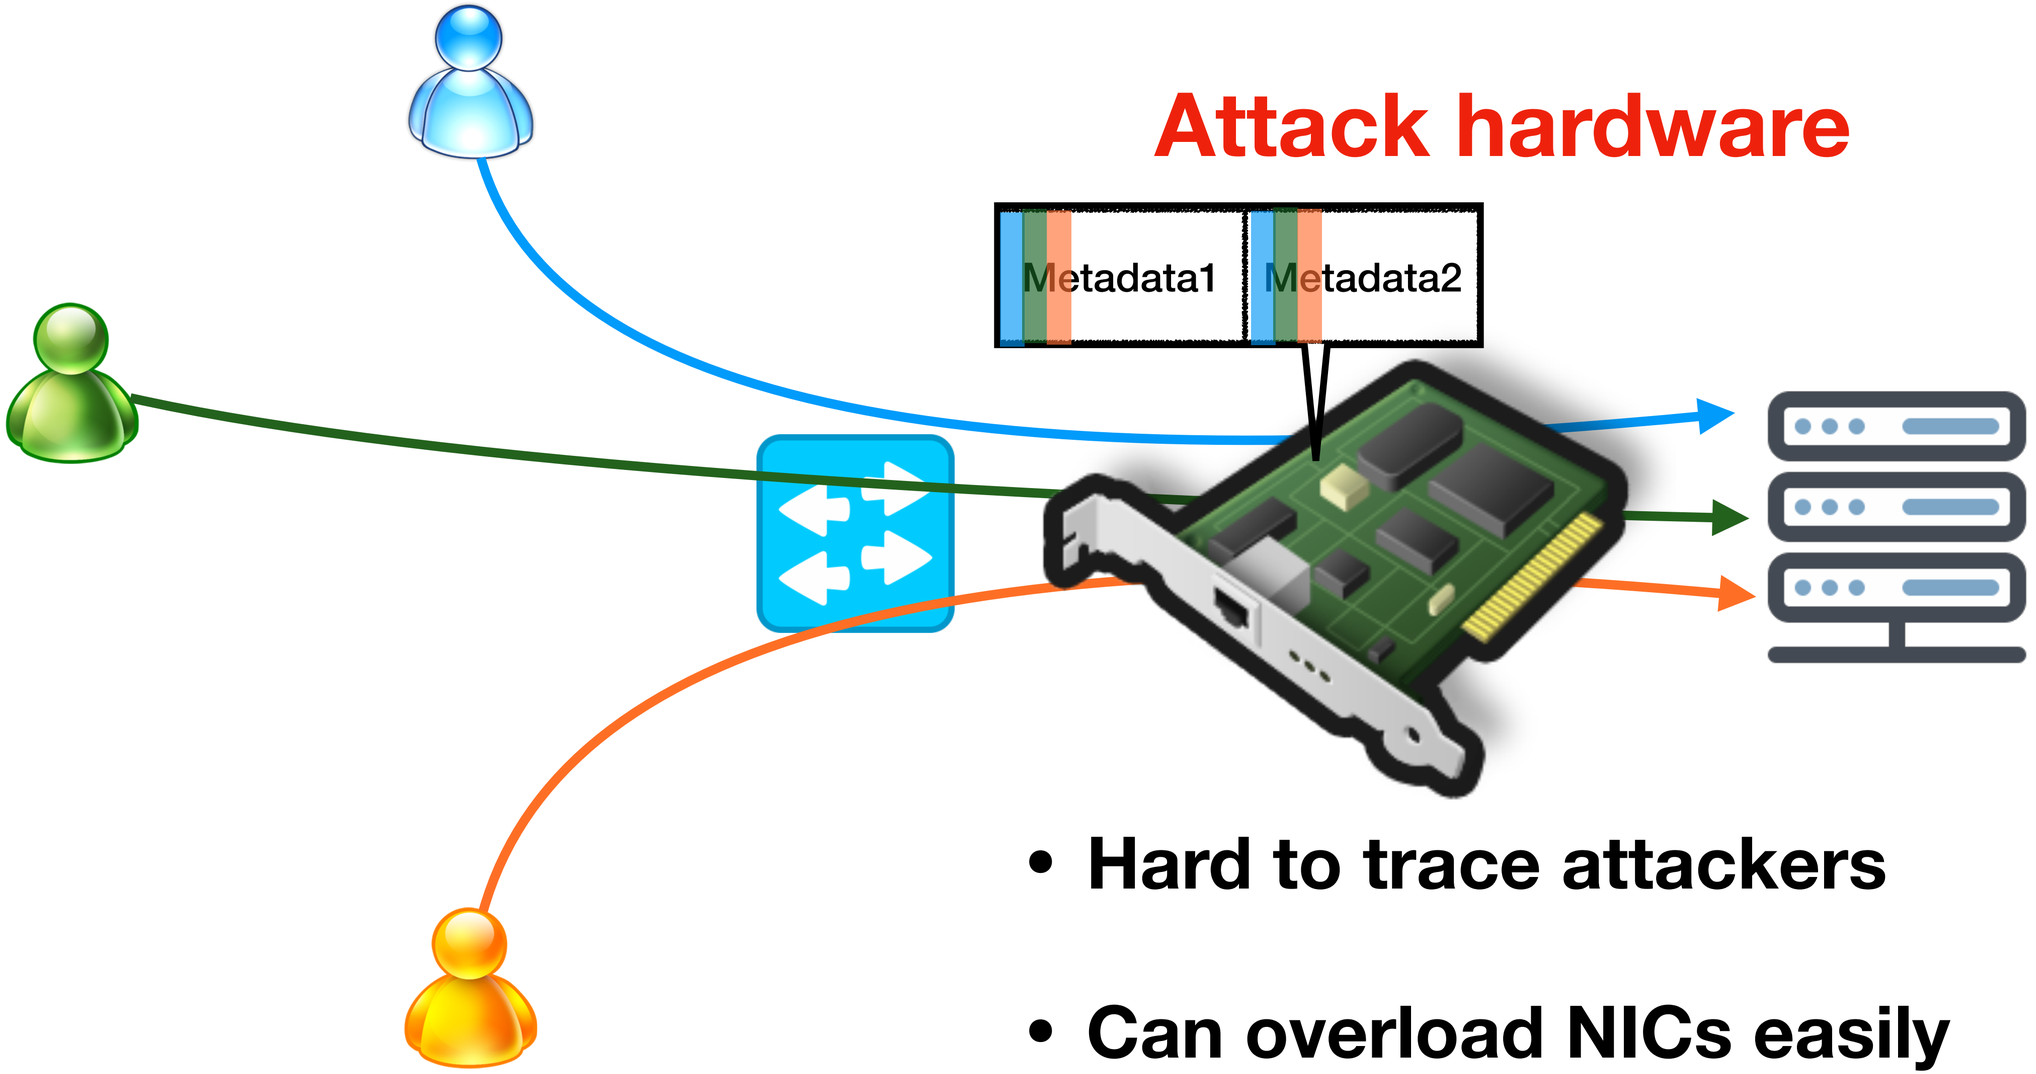

- **Hard to trace attackers**
- **Can overload NICs easily**

Vulnerability 2: Denial of Service

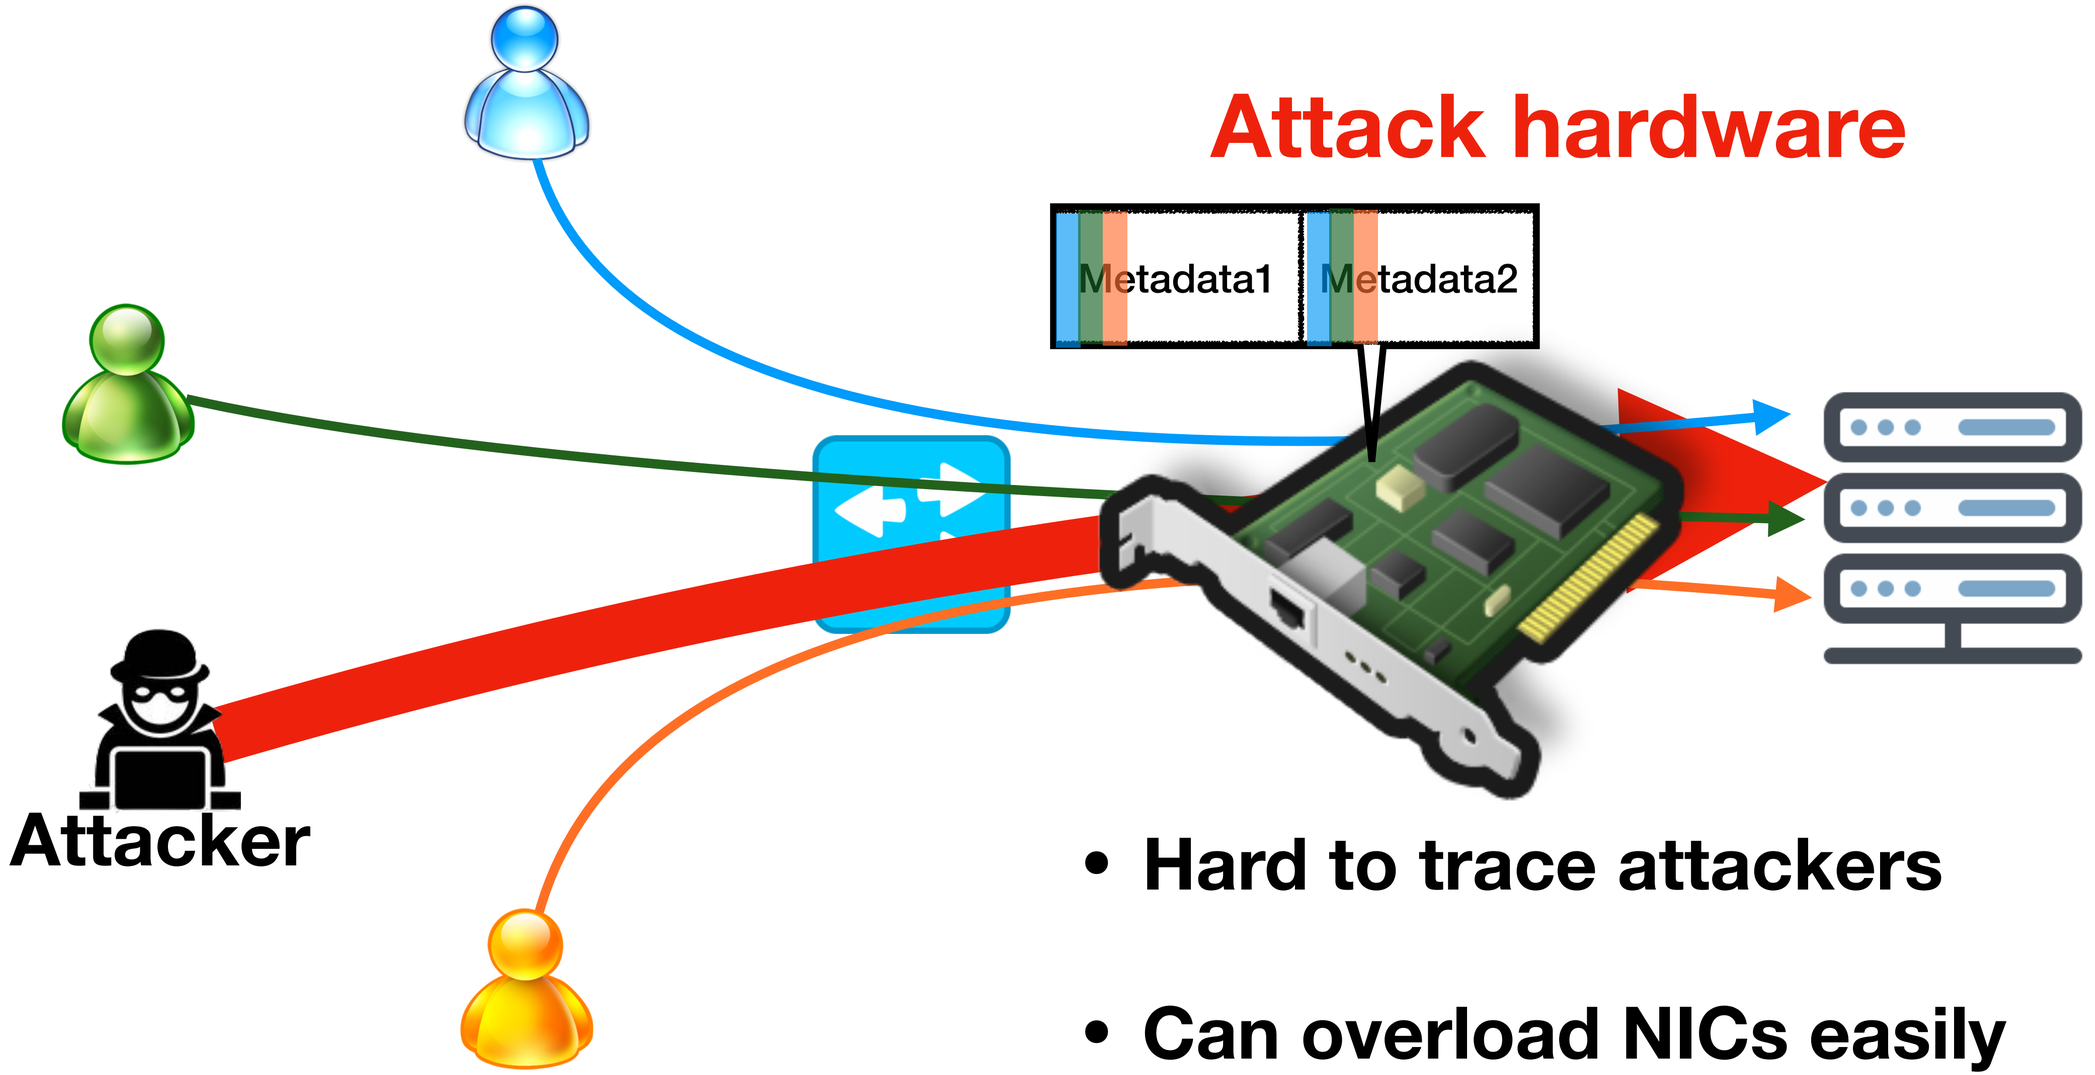

- **Hard to trace attackers**
- **Can overload NICs easily**

Vulnerability 2: Denial of Service

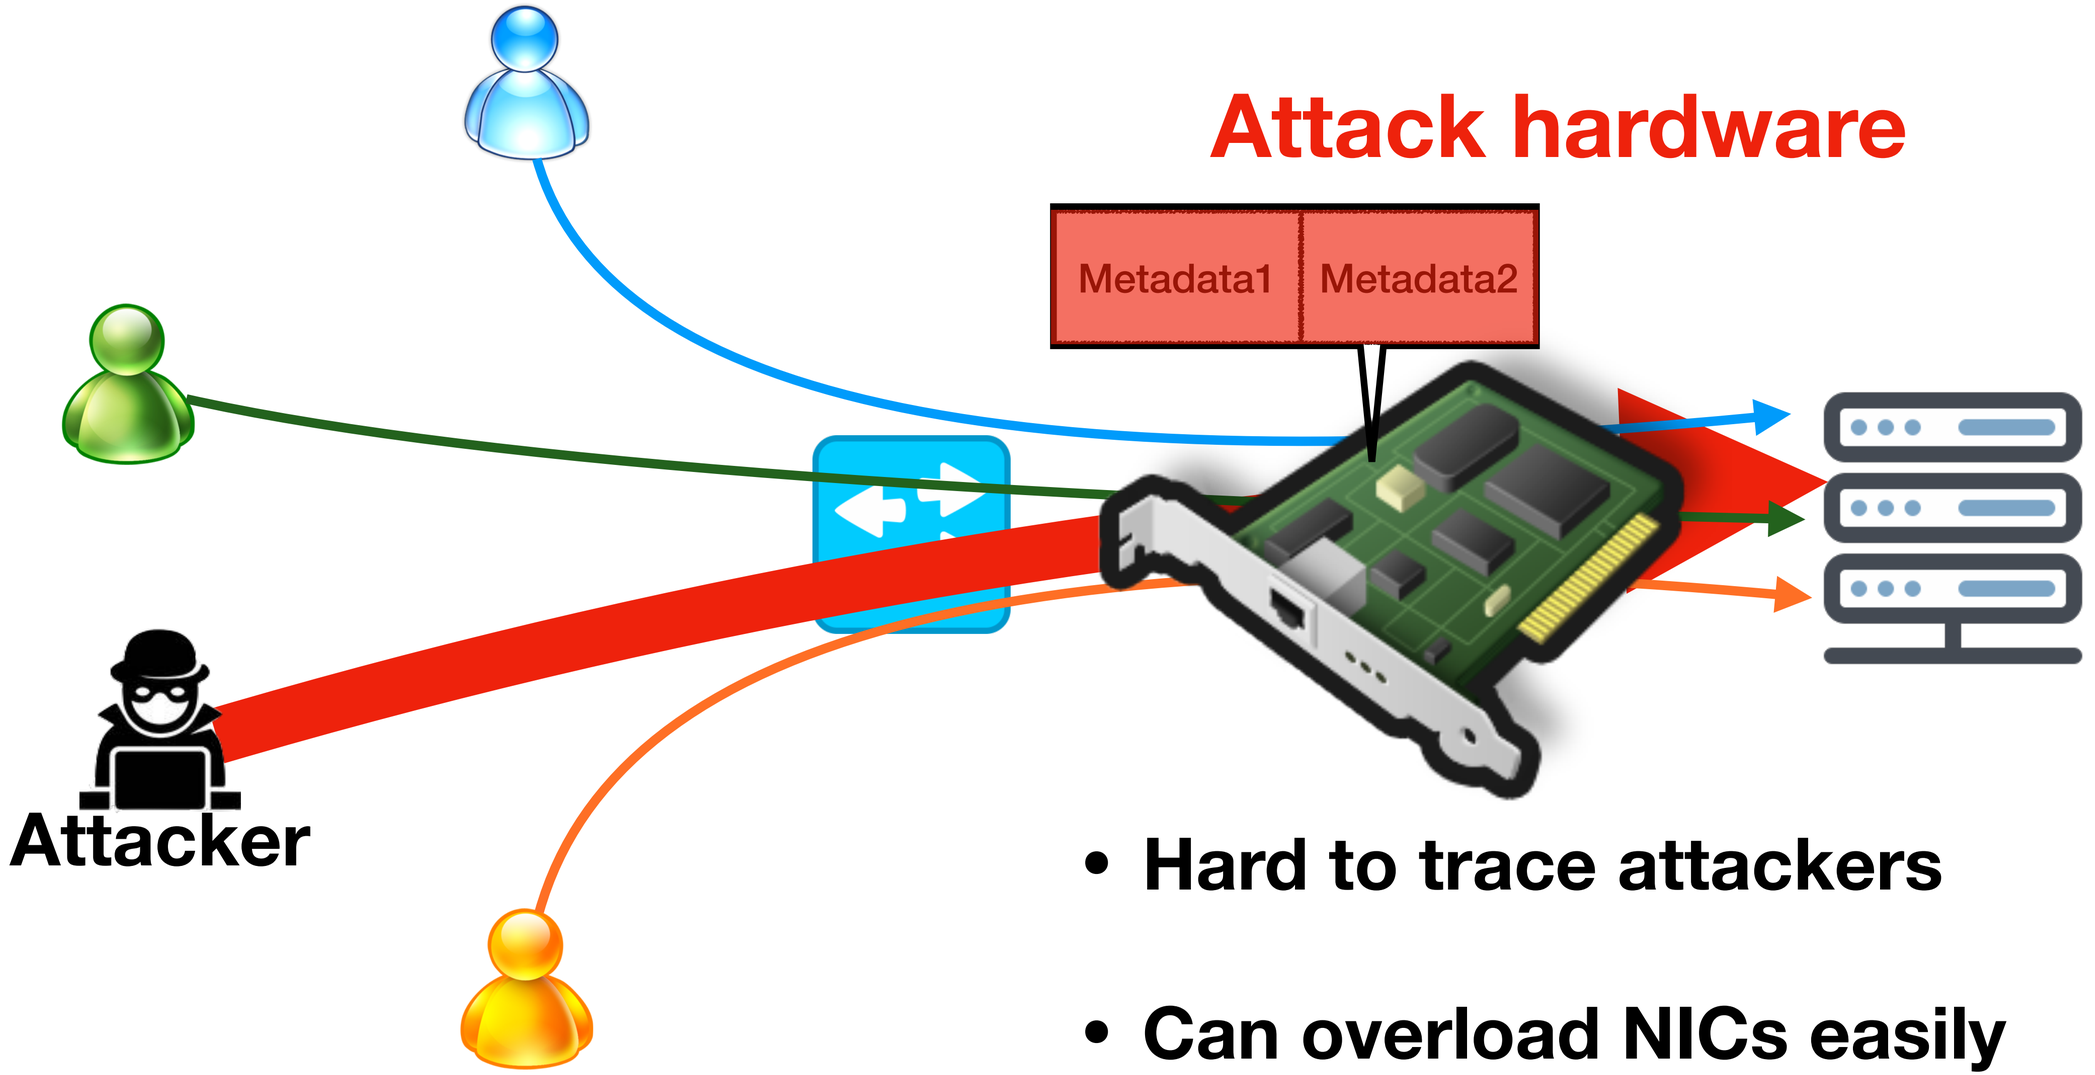

- **Hard to trace attackers**
- **Can overload NICs easily**

Vulnerability 2: Denial of Service

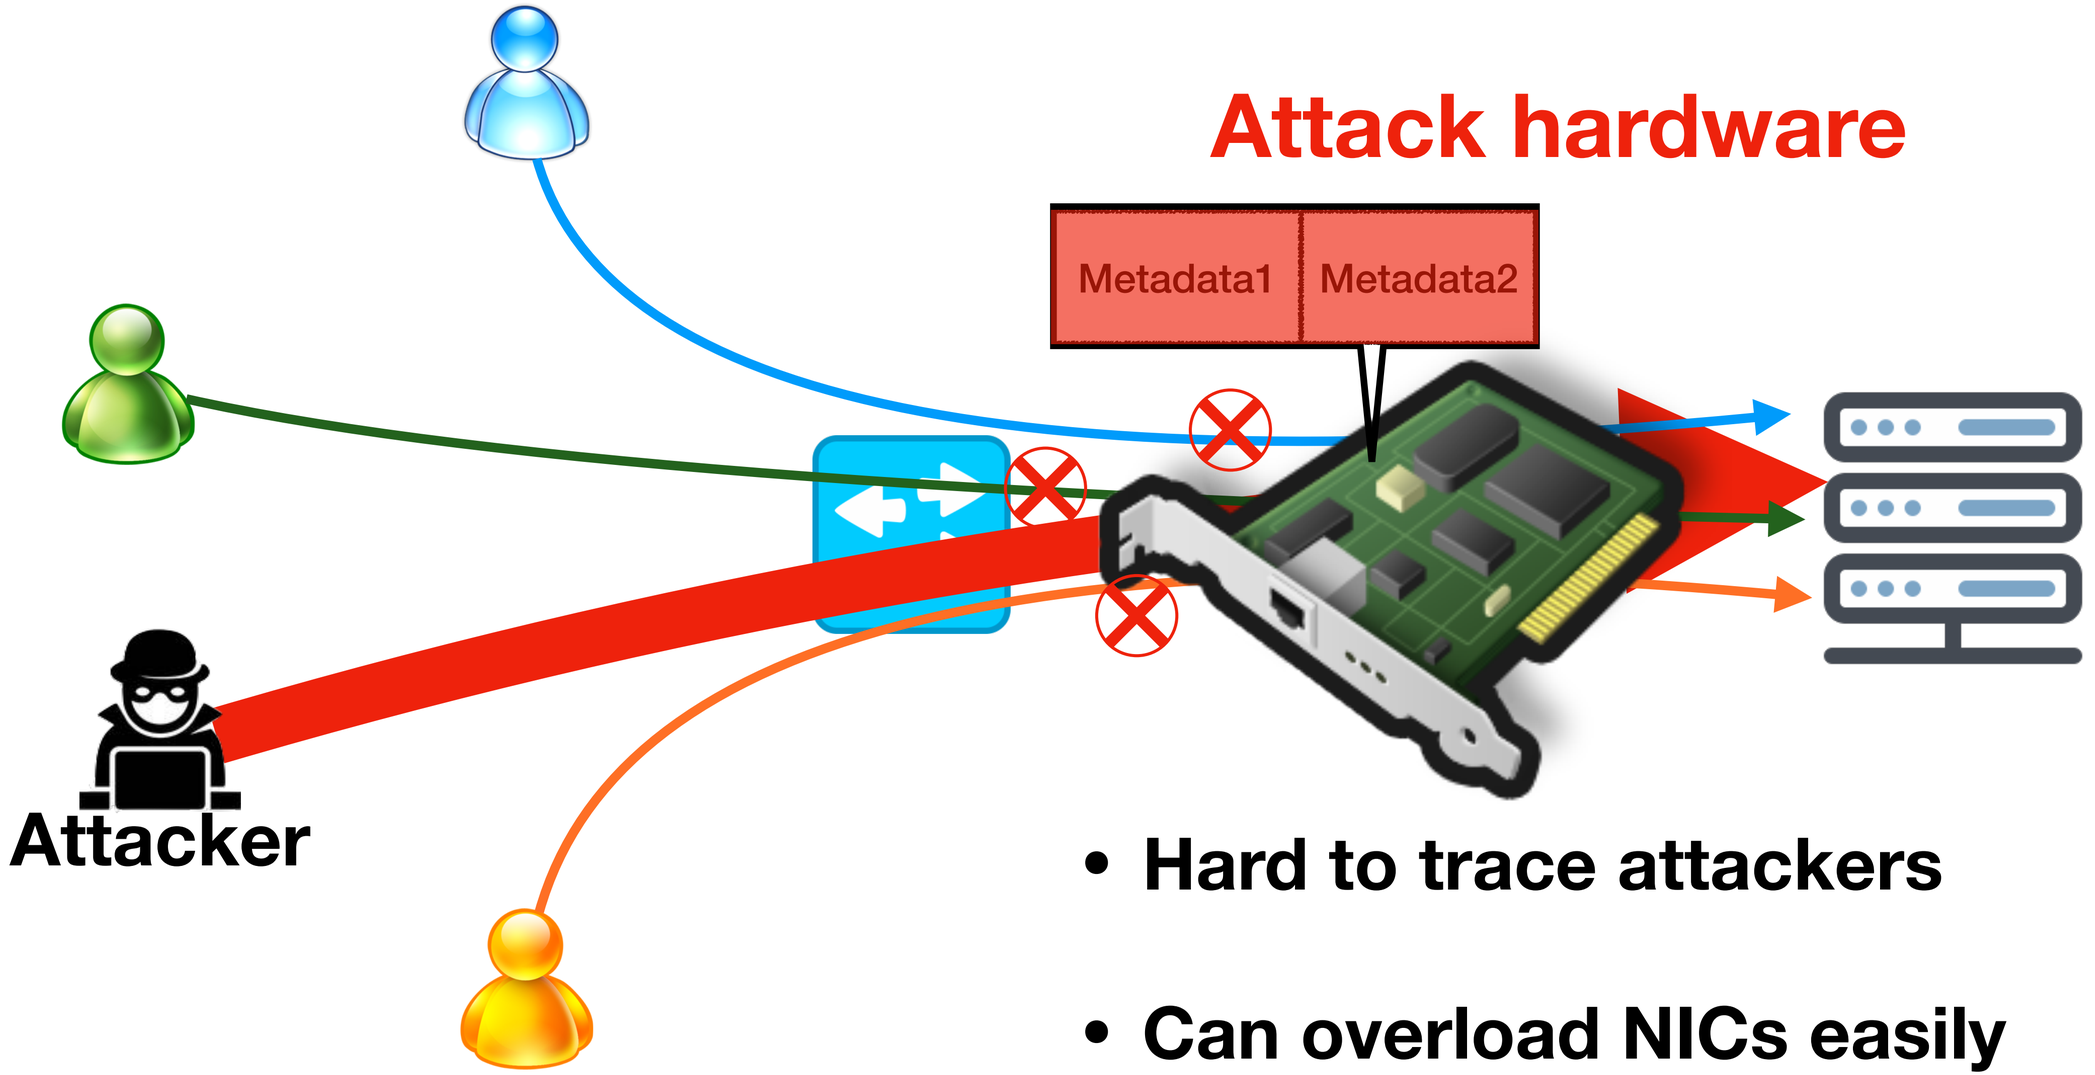

- **Hard to trace attackers**
- **Can overload NICs easily**

Discussion and Defense

- Adding intermediate layer at the sender side

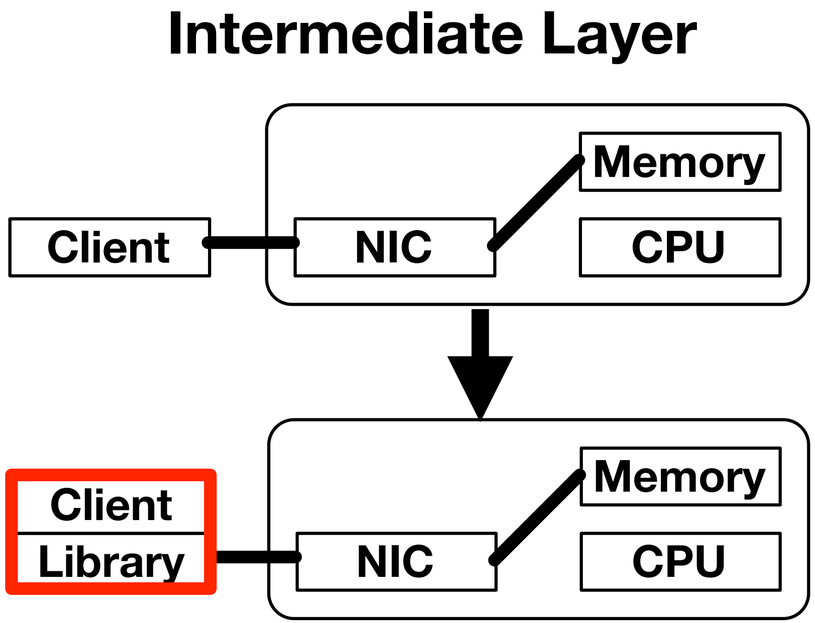

Discussion and Defense

- Adding intermediate layer at the sender side
- Enhancing SmartNIC at the receiver side

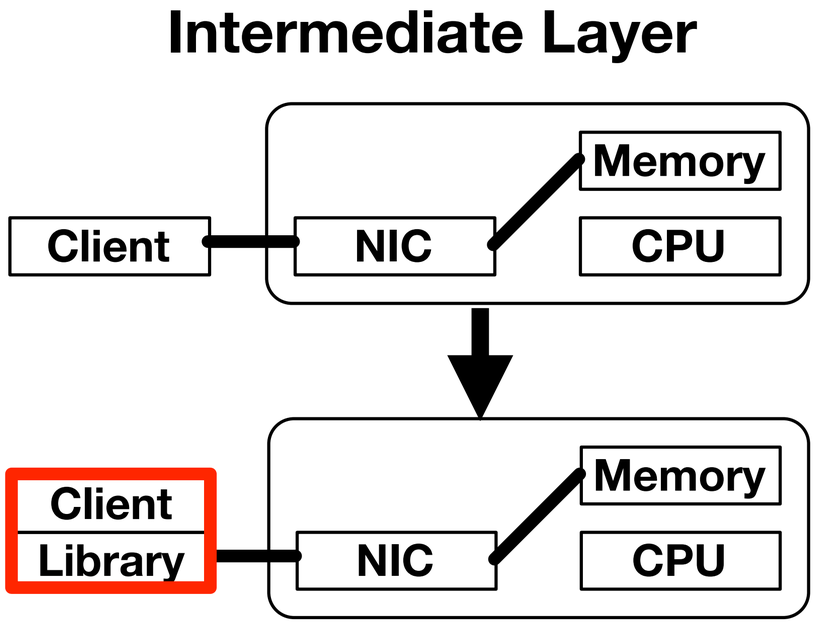

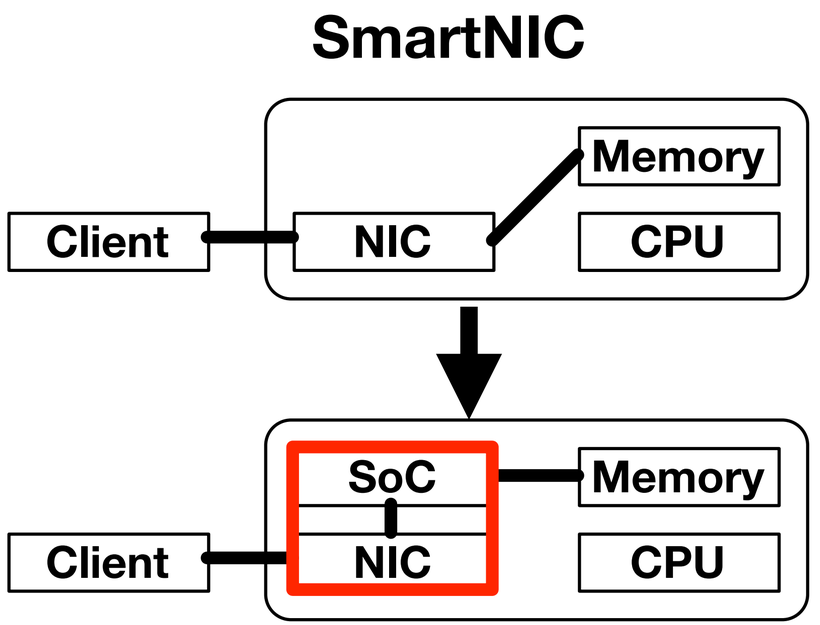

Outline

- Introduction and Background
- Vulnerabilities in One-Sided Communication
- Vulnerabilities in One-Sided Hardware
- Opportunities in One-Sided Communication
- Conclusion

One- and Two-Sided Hardware

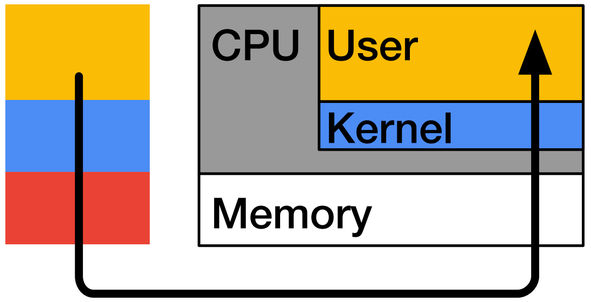

Two-Sided

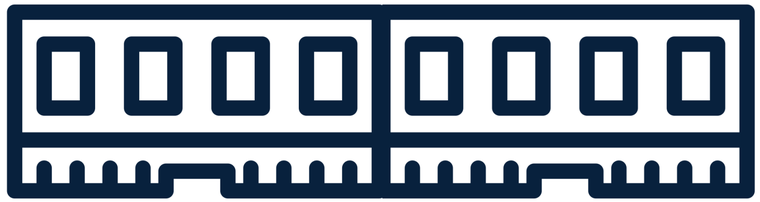

One- and Two-Sided Hardware

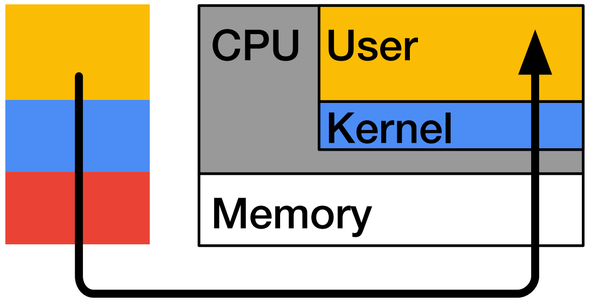

Two-Sided

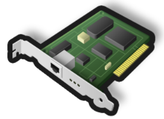

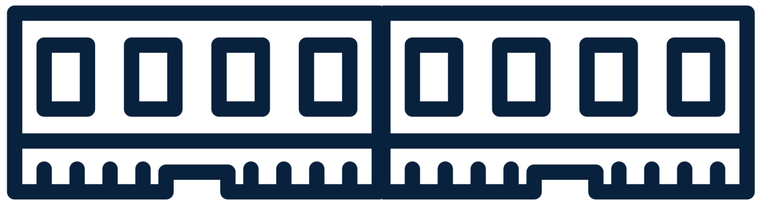

One- and Two-Sided Hardware

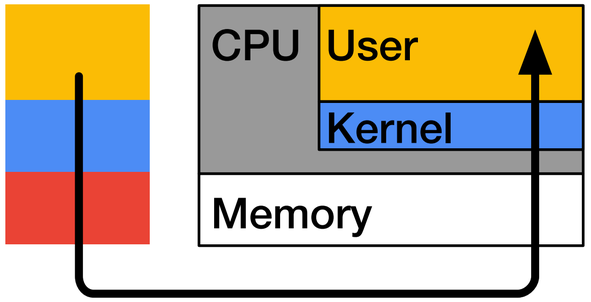

Two-Sided

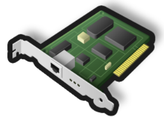

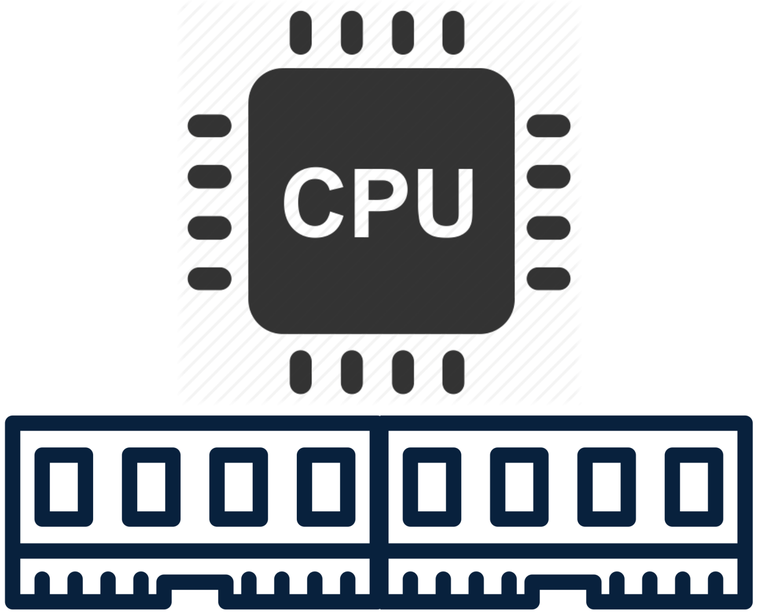

One- and Two-Sided Hardware

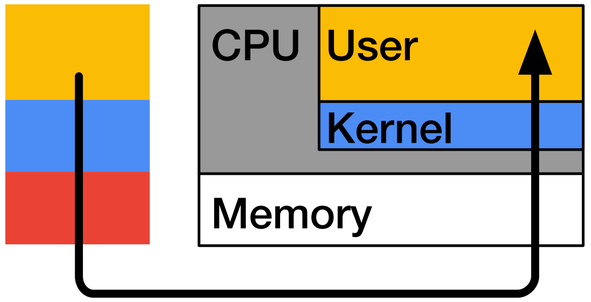

Two-Sided

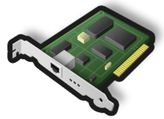

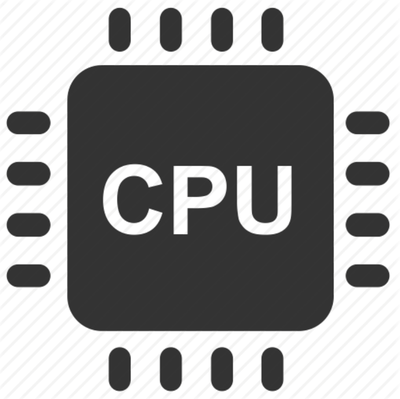

1. Address mapping
2. Permission checking
3. Resource isolation

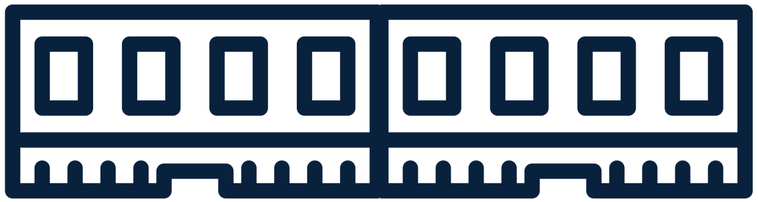

One- and Two-Sided Hardware

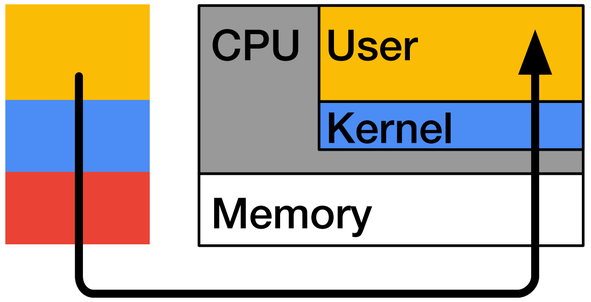

Two-Sided

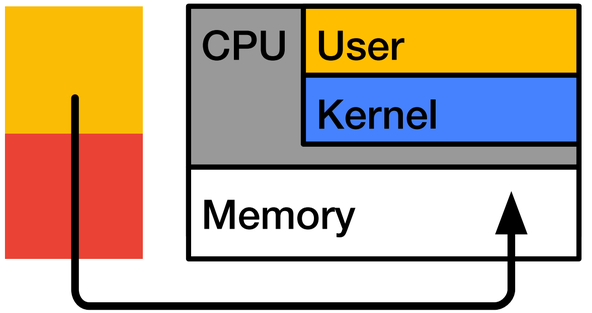

One-Sided

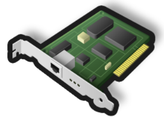

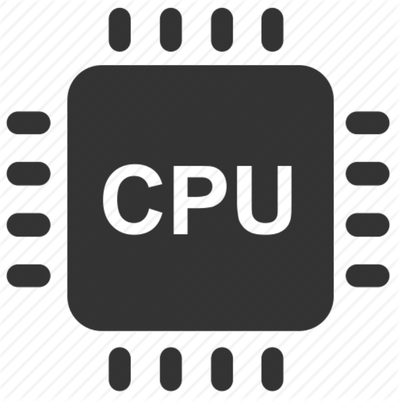

1. Address mapping
2. Permission checking
3. Resource isolation

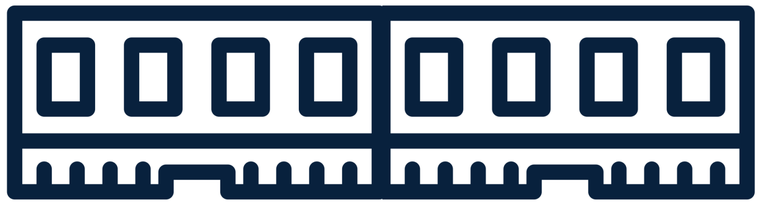

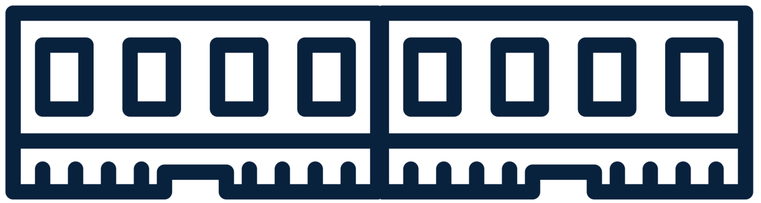

One- and Two-Sided Hardware

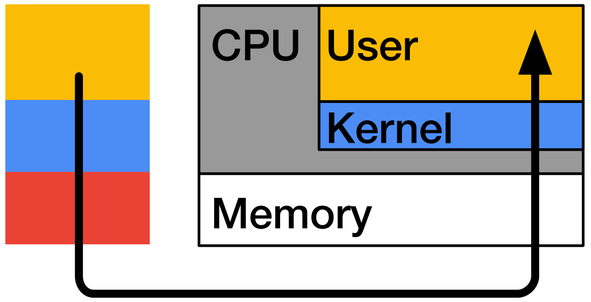

Two-Sided

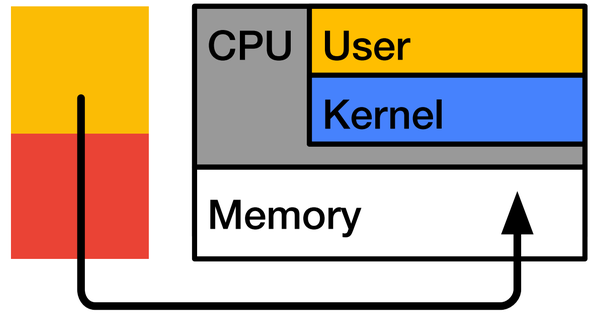

One-Sided

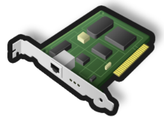

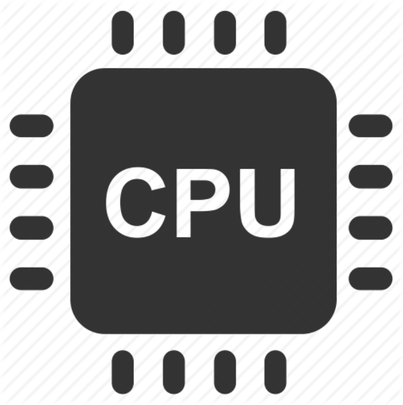

1. Address mapping
2. Permission checking
3. Resource isolation

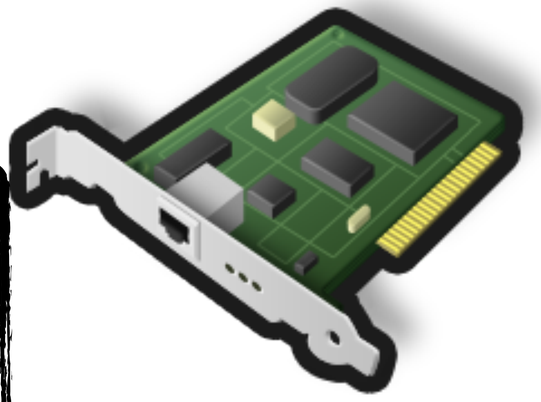

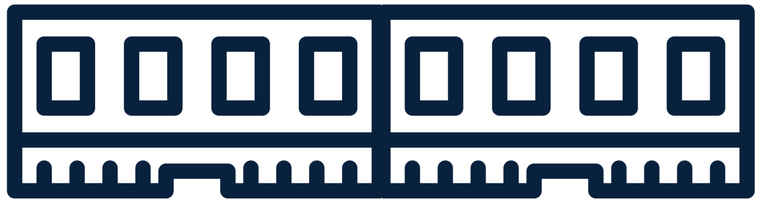

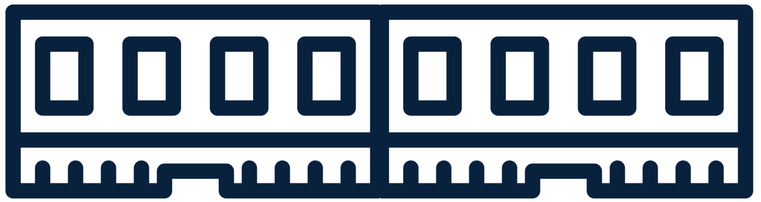

One- and Two-Sided Hardware

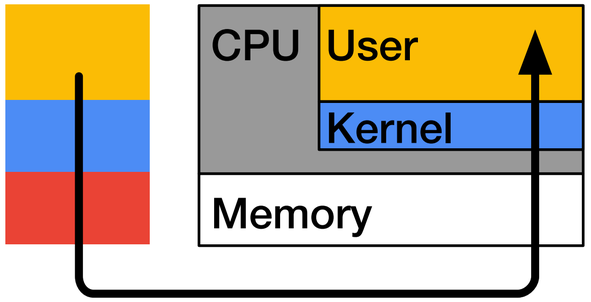

Two-Sided

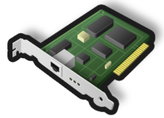

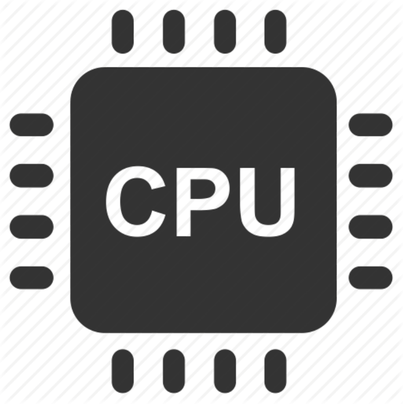

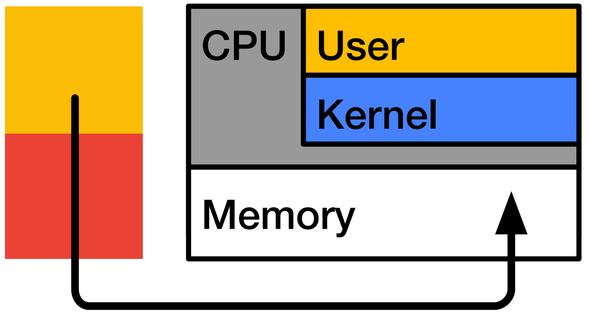

One-Sided

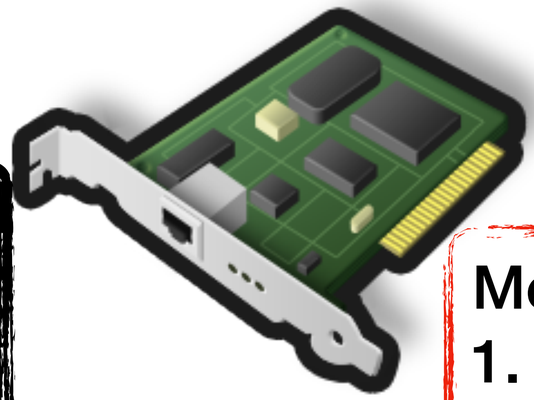

1. Address mapping
2. Permission checking
3. Resource isolation

Memory Region
1. rkey/lkey
2. Address

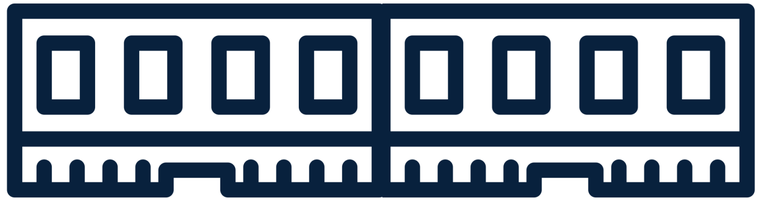

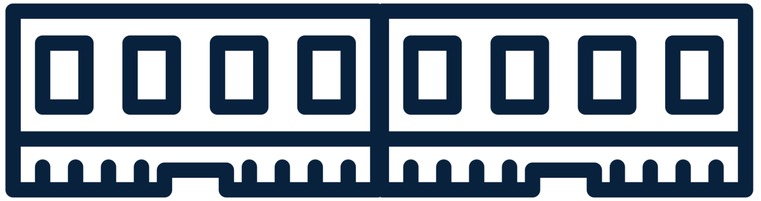

Vulnerability 3 - Predictable Hardware Managed Keys

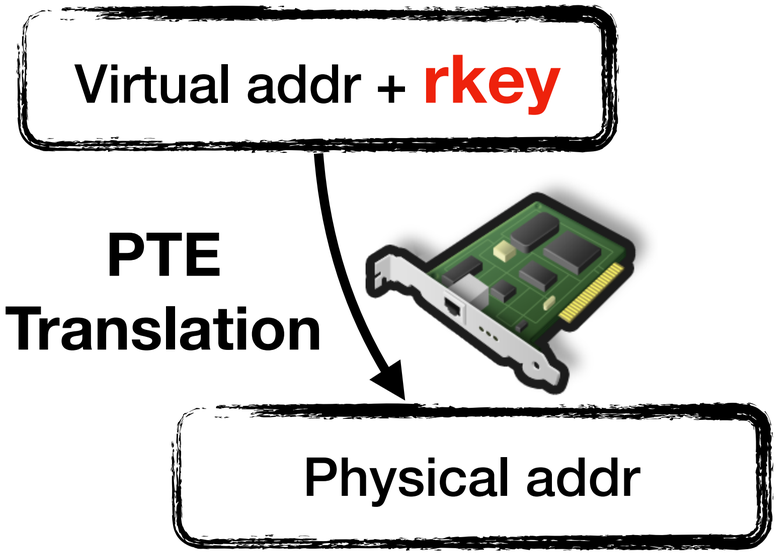

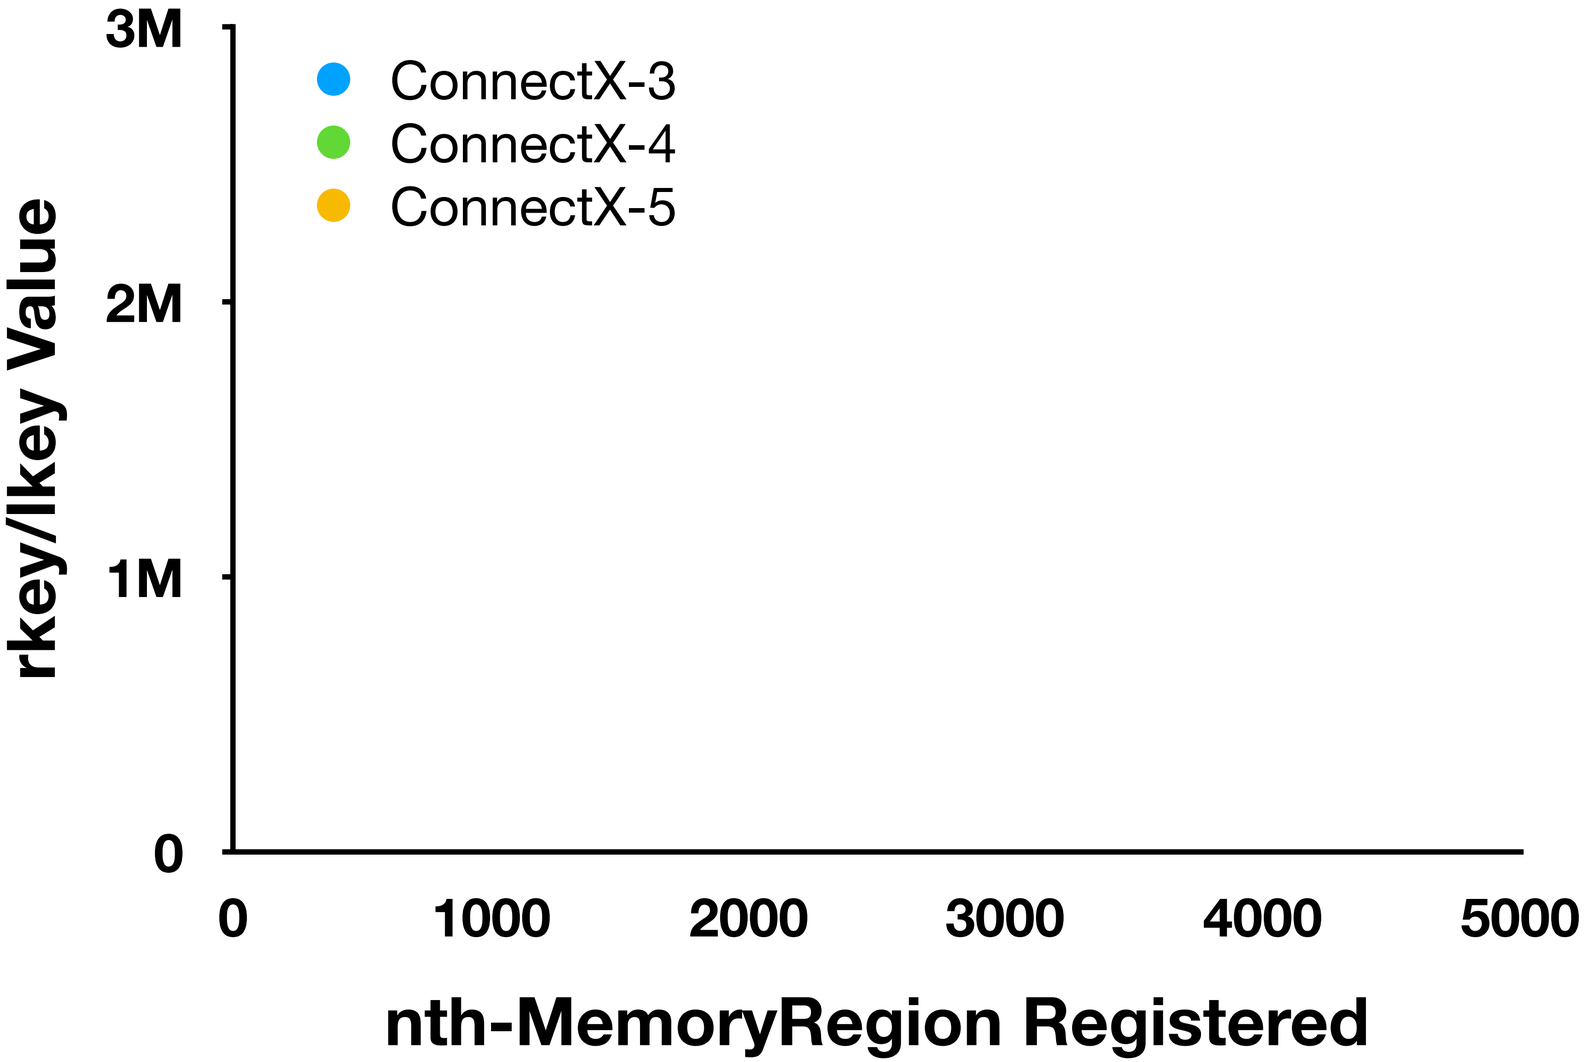

Vulnerability 3 - Predictable Hardware Managed Keys

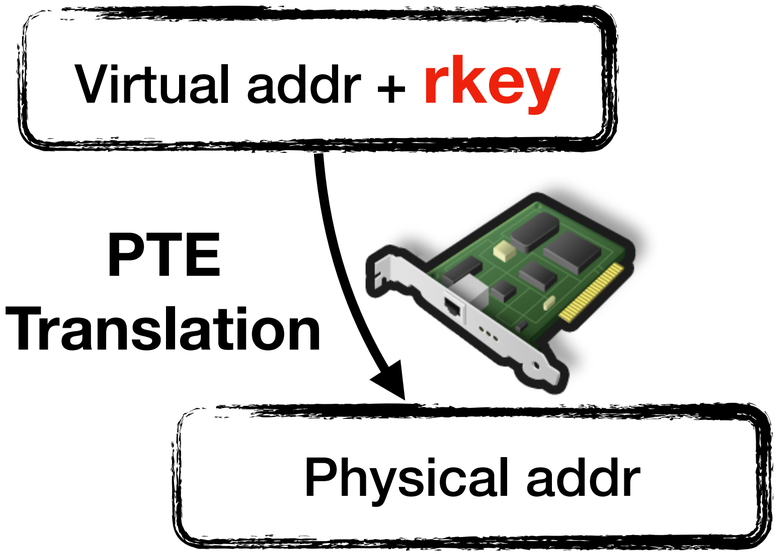

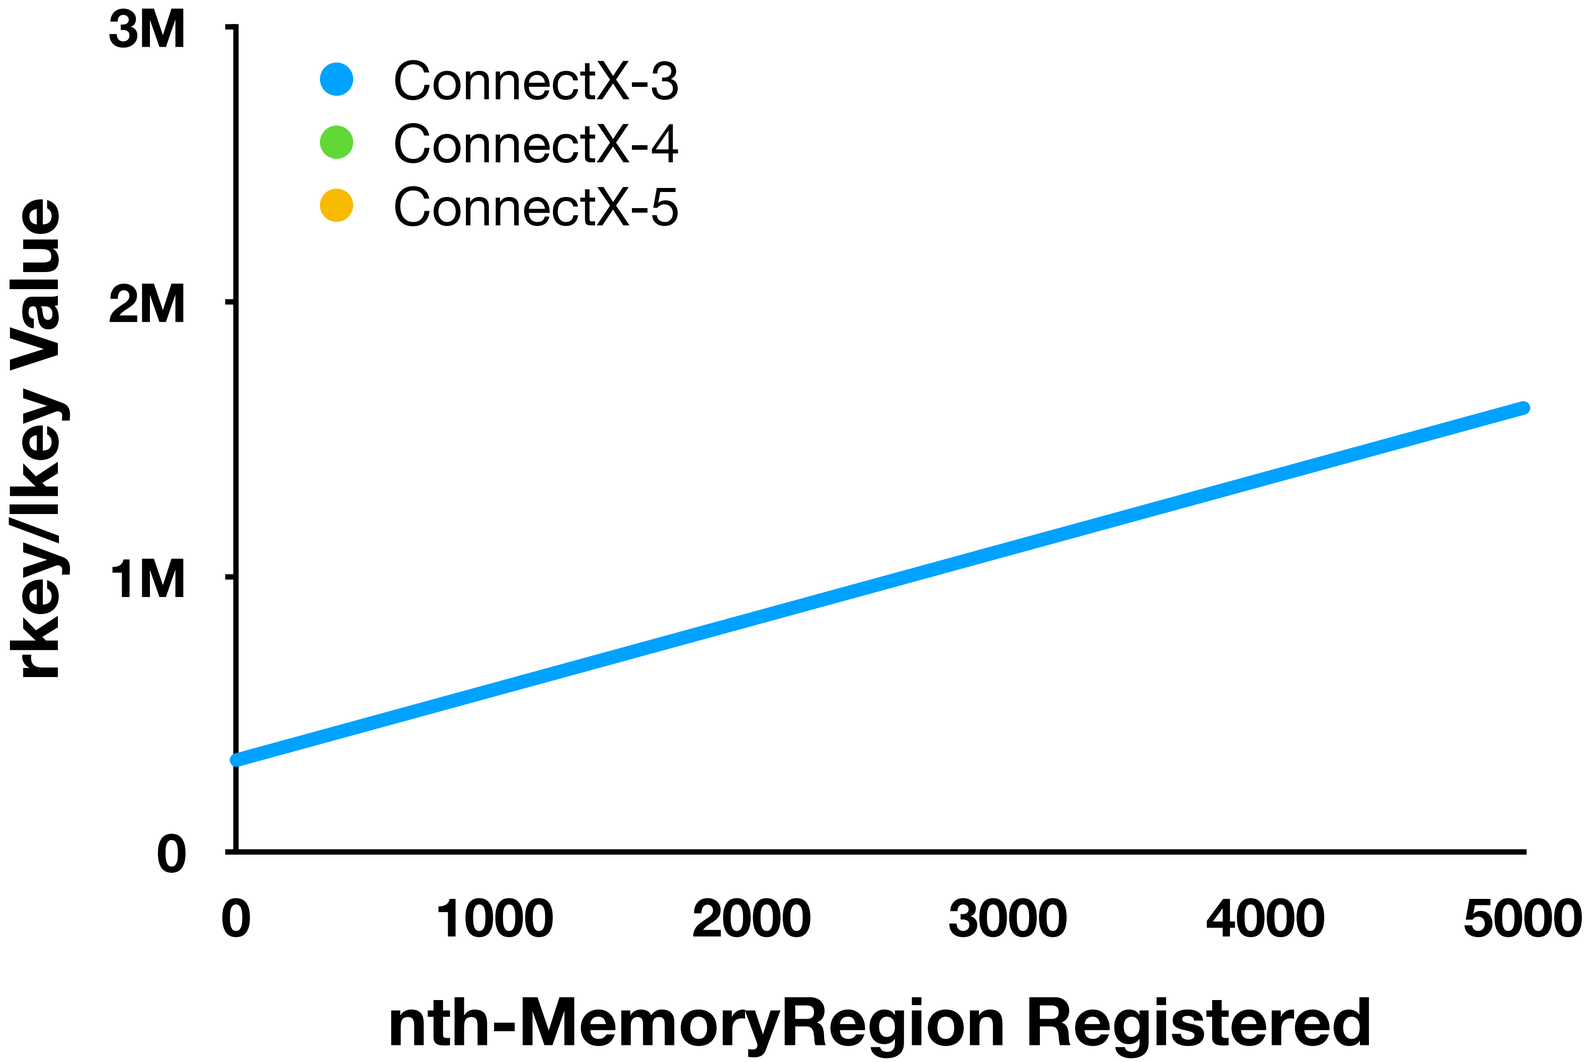

Vulnerability 3 - Predictable Hardware Managed Keys

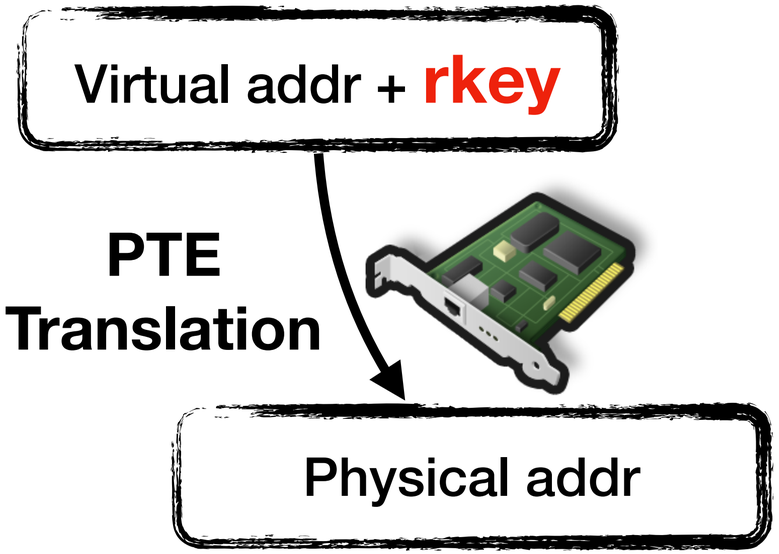

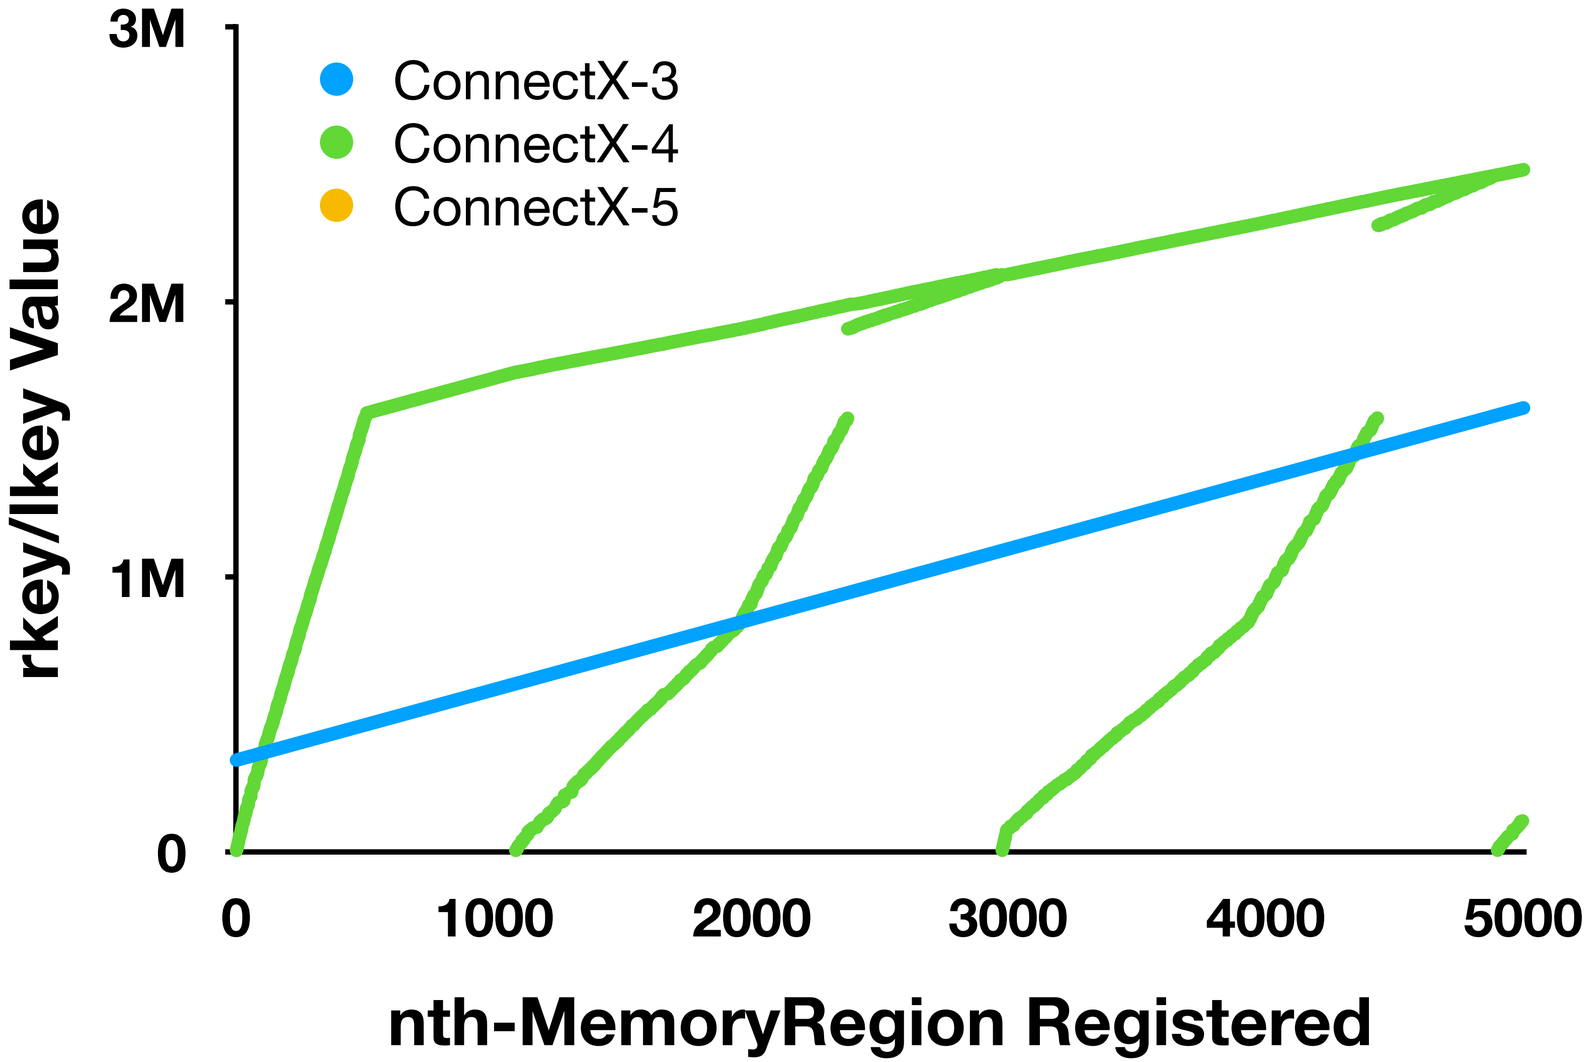

Vulnerability 3 - Predictable Hardware Managed Keys

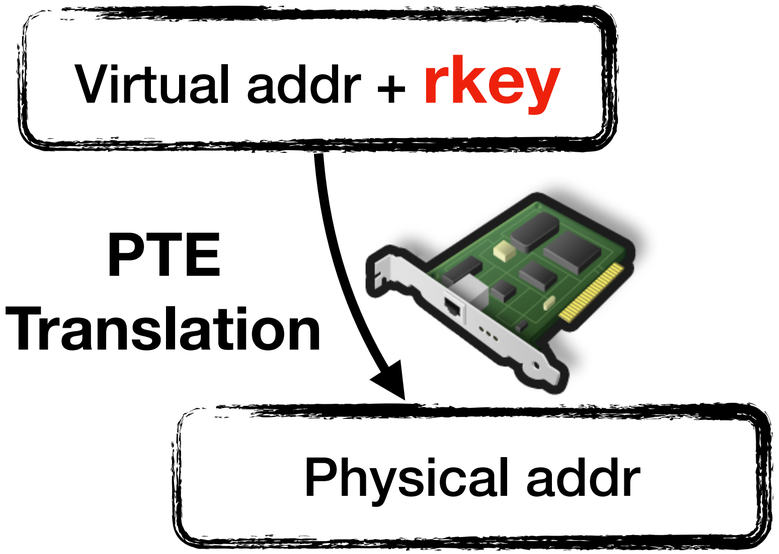

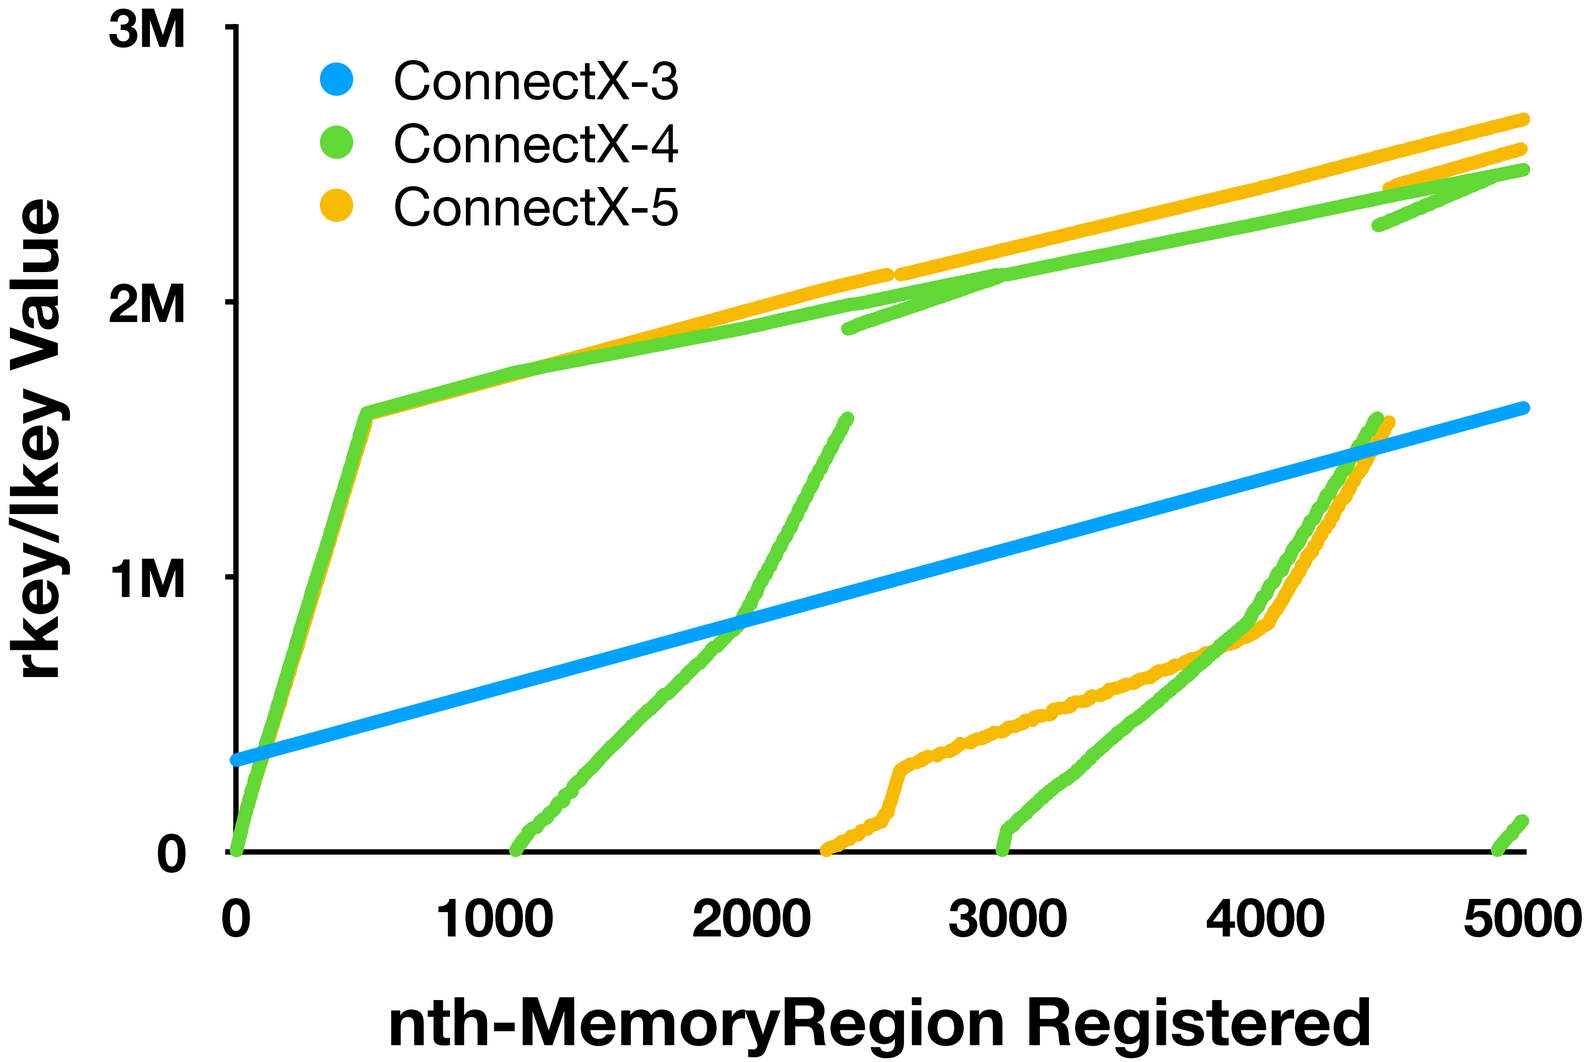

Vulnerability 4 - Side Channel in NICs

ConnectX-5, 1KB READ request latency

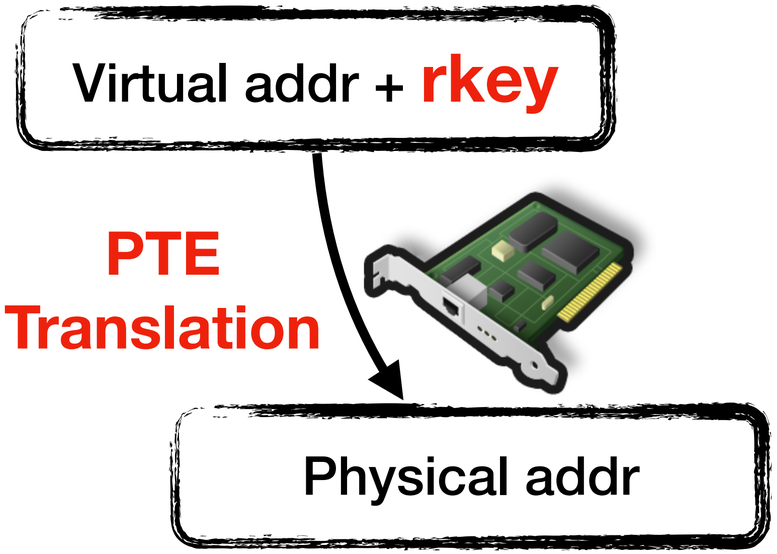

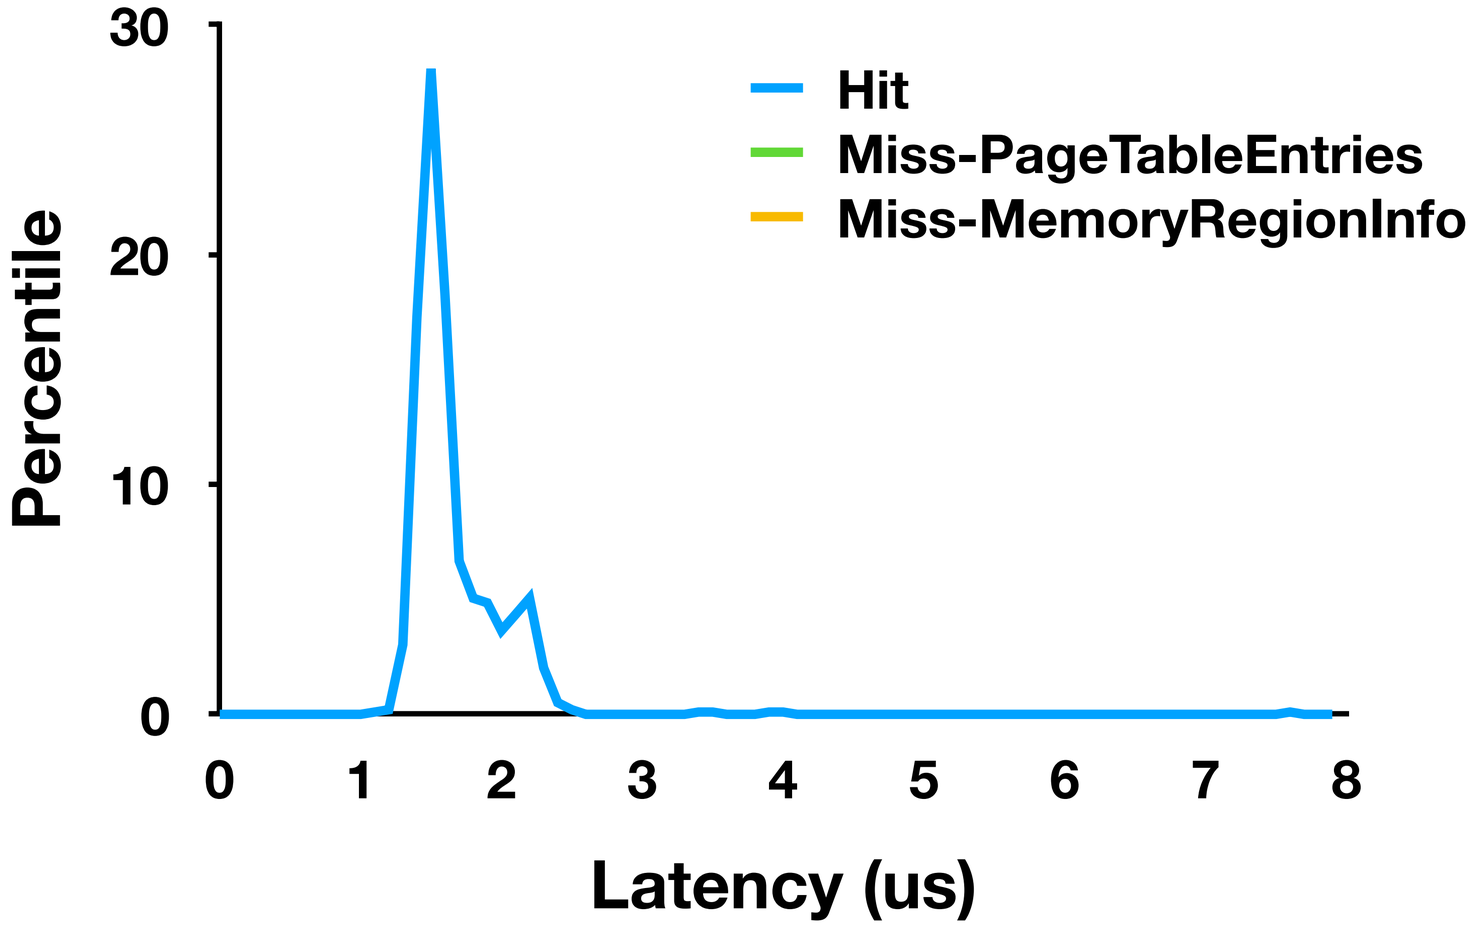

Vulnerability 4 - Side Channel in NICs

ConnectX-5, 1KB READ request latency

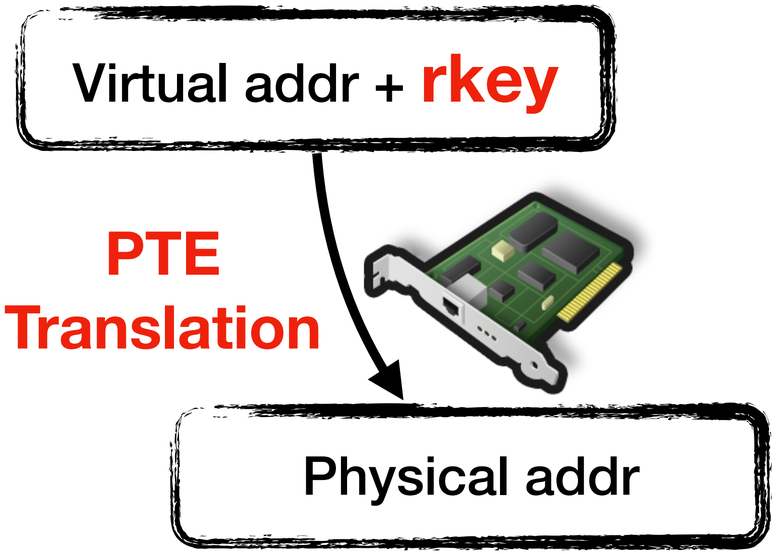

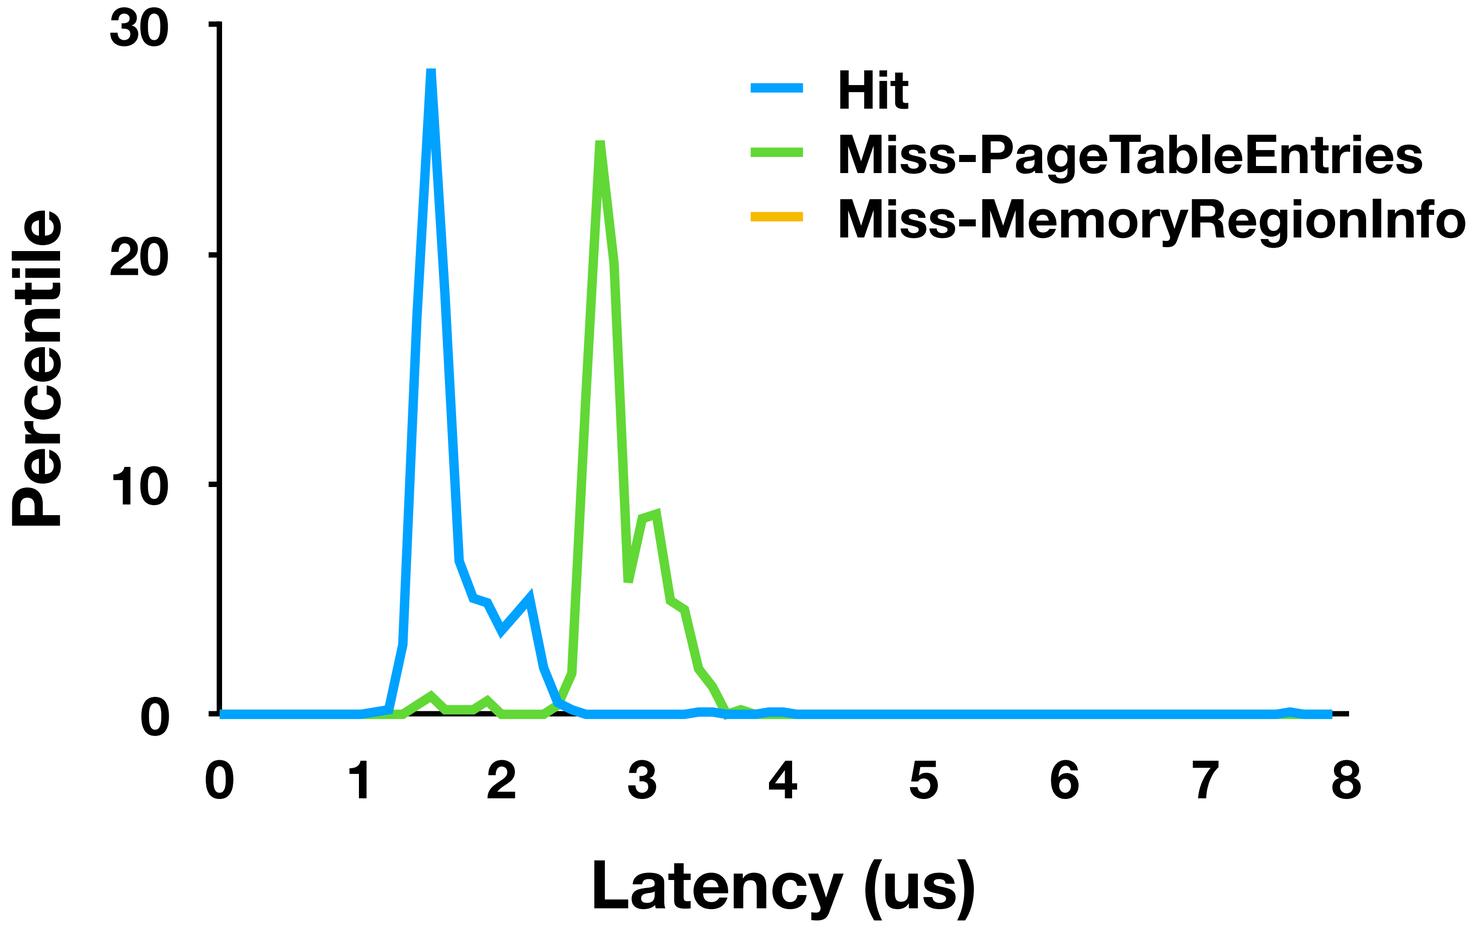

Vulnerability 4 - Side Channel in NICs

ConnectX-5, 1KB READ request latency

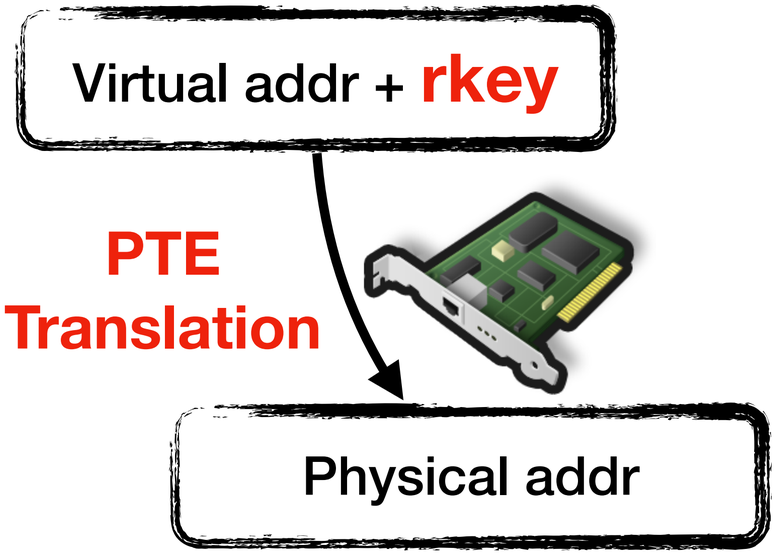

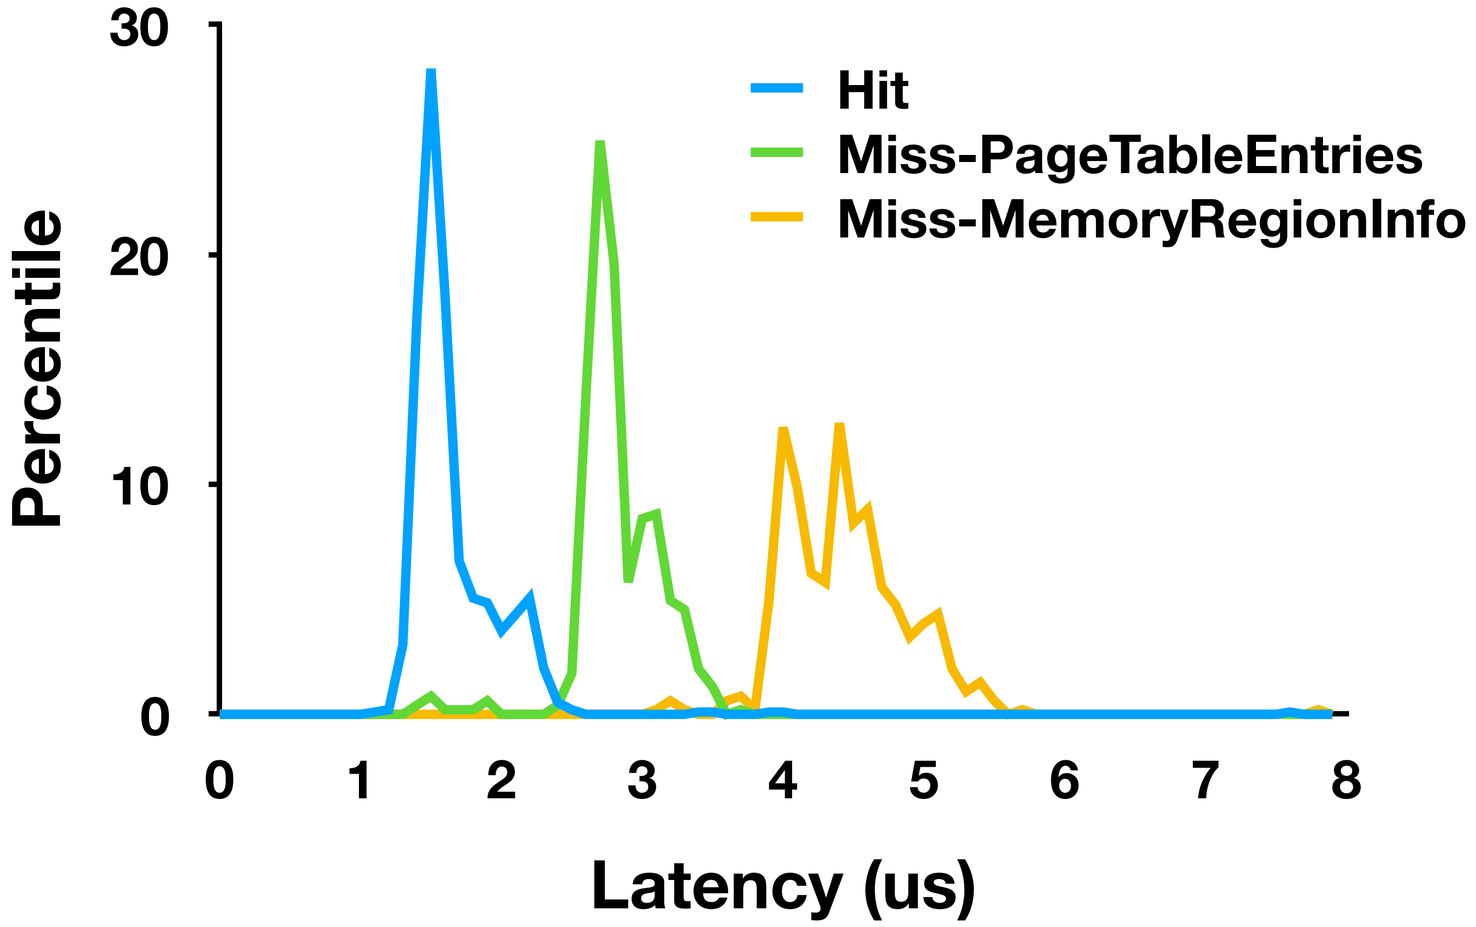

Side-Channel Attacks in RDMA (Pythia, USENIX Sec '19)

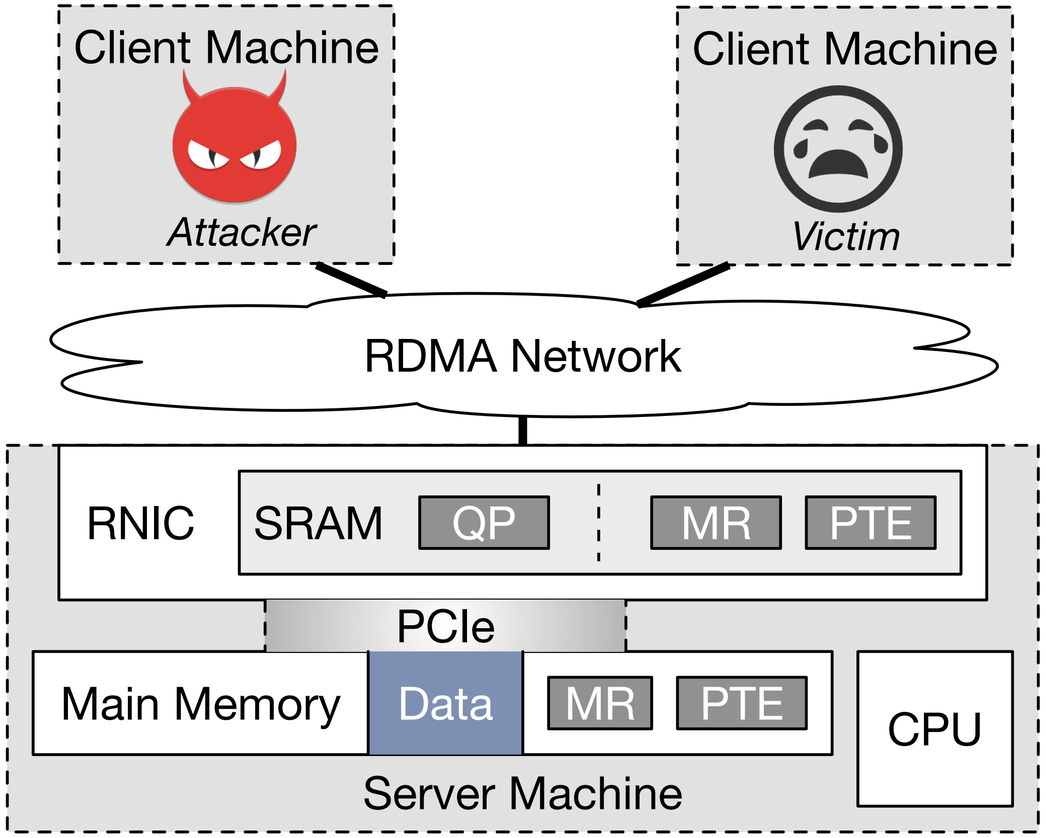

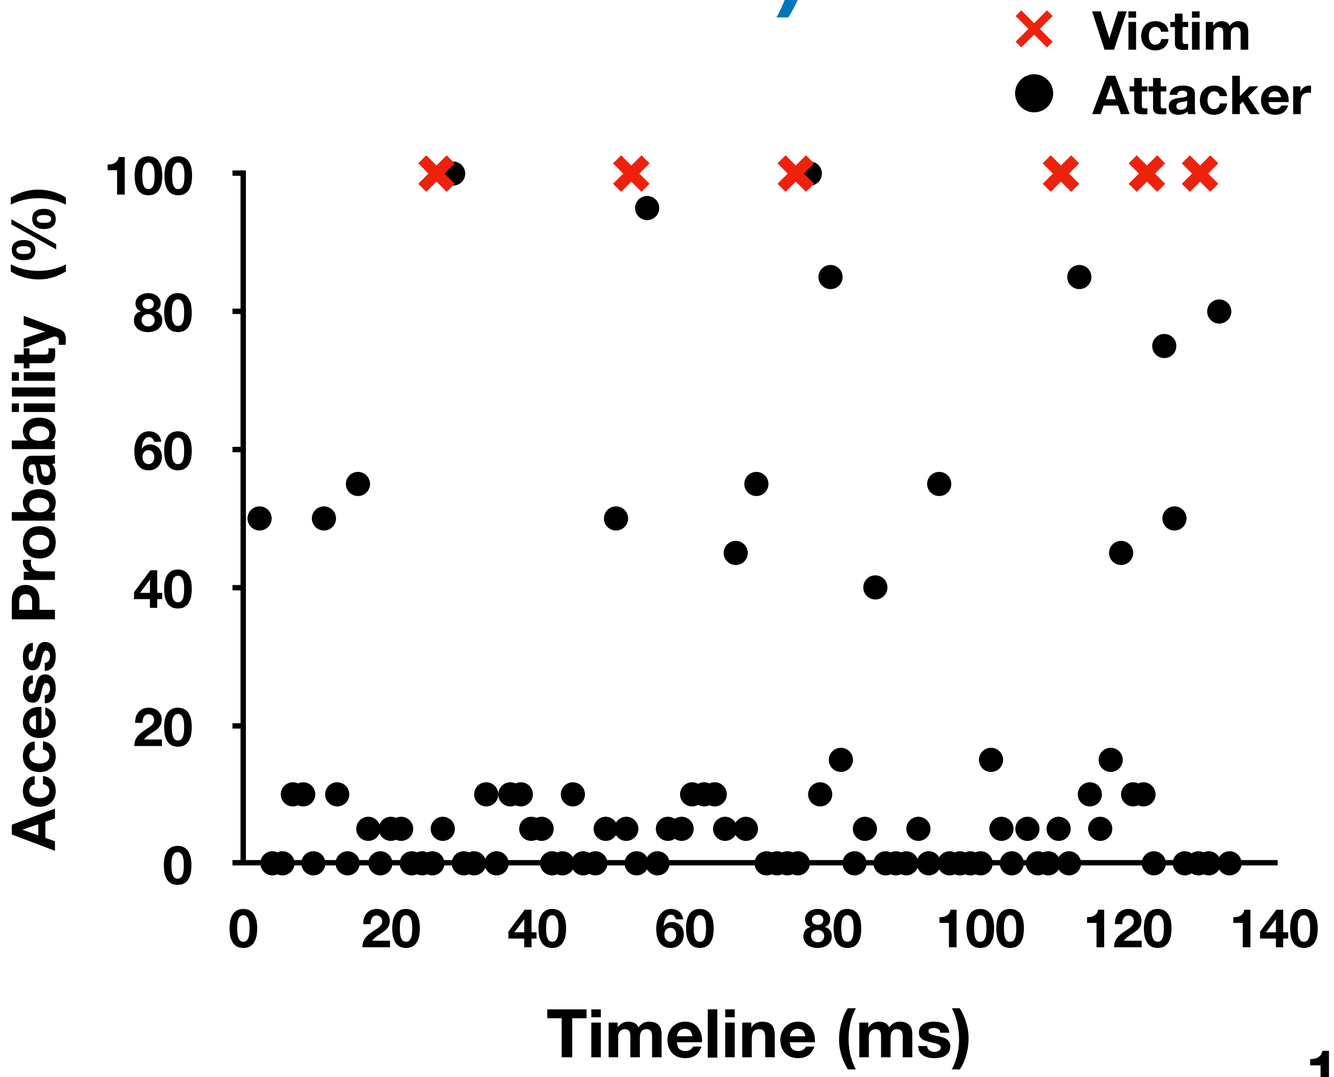

Discussion and Defense

- Generate memory registration keys cryptographically

Sequential to Random

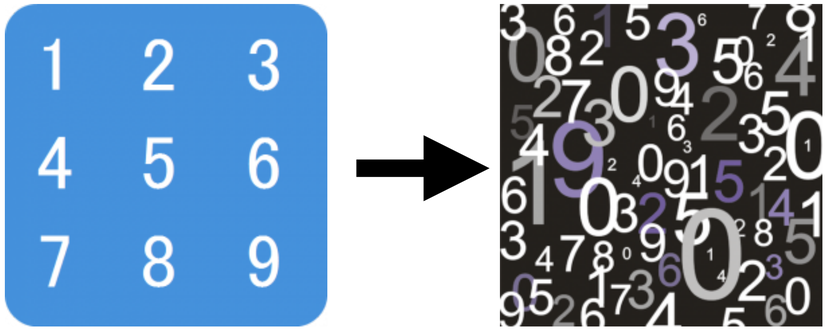

Discussion and Defense

- Generate memory registration keys cryptographically
- Isolate on-board resources for different clients

Sequential to Random

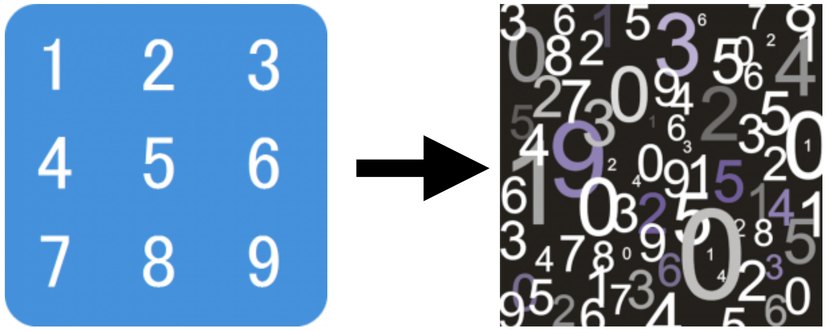

On-board Resource

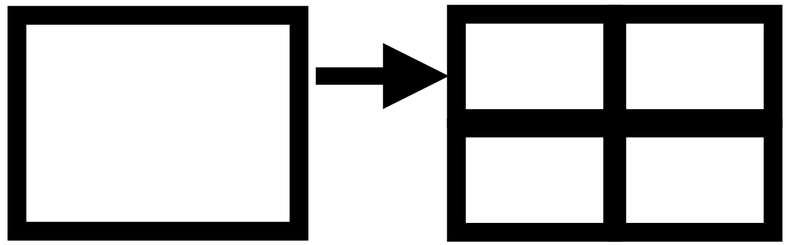

Discussion and Defense

- Generate memory registration keys cryptographically
- Isolate on-board resources for different clients
- Enhancing SmartNIC at the receiver side

Sequential to Random

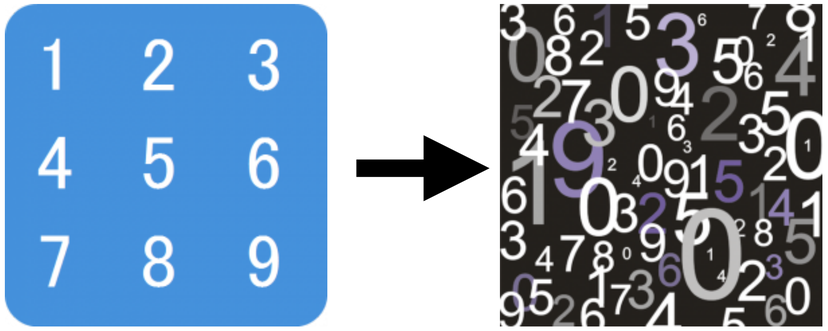

On-board Resource

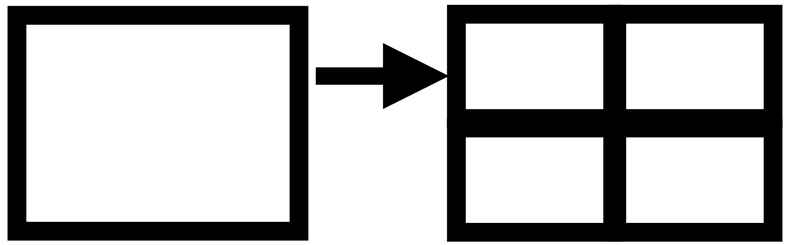

SmartNIC

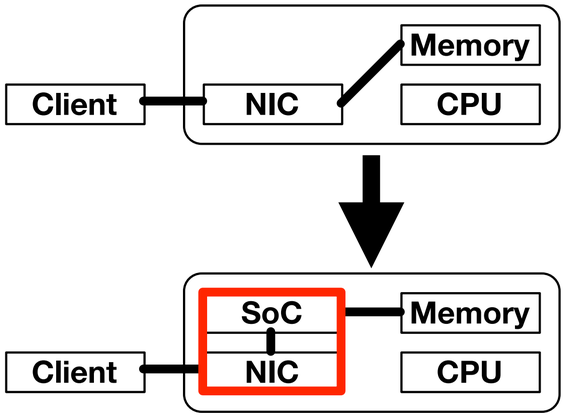

Outline

- Introduction and Background
- Vulnerabilities in One-Sided Communication
- Vulnerabilities in One-Sided Hardware
- Opportunities in One-Sided Communication
- Conclusion

Opportunity of One-sided Communication

ORAM Access

Server

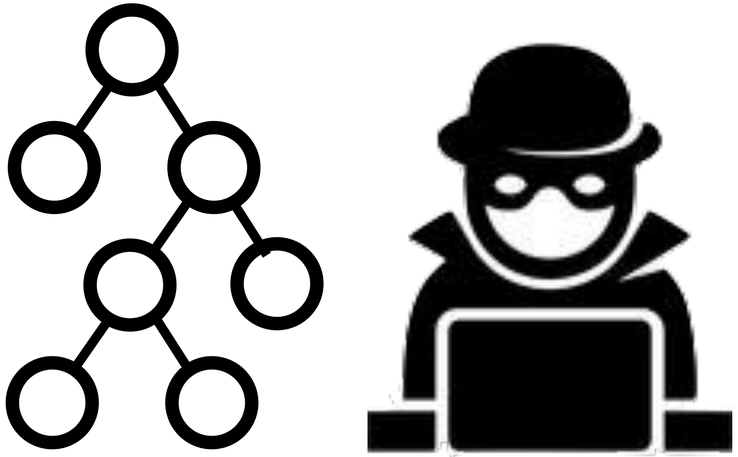

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

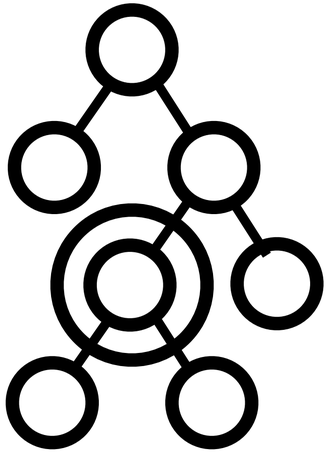

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

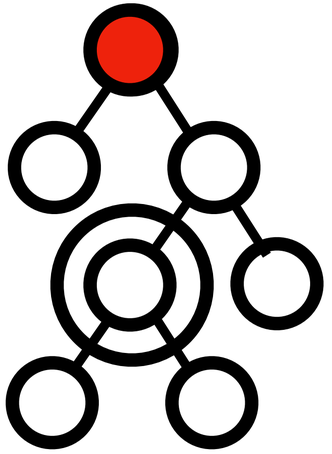

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

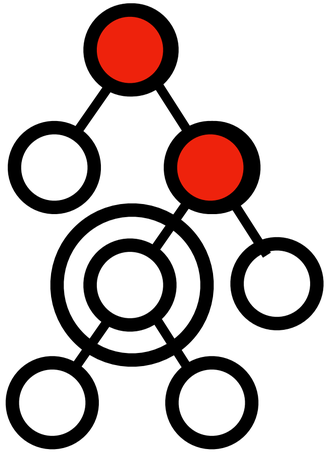

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

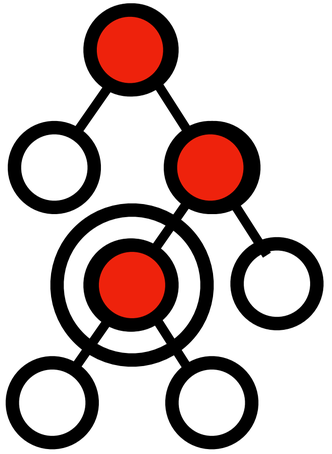

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

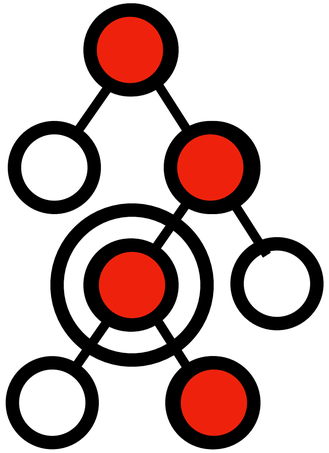

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

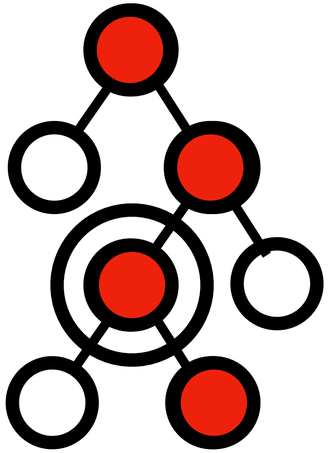

One-sided READ

Opportunity of One-sided Communication

ORAM Access

Server

ORAM READ/WRITE

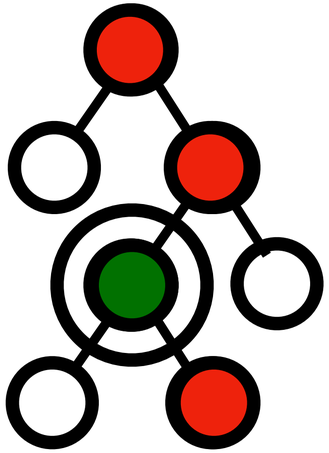

One-sided READ

Opportunity of One-sided Communication

ORAM Access

Server

(1-K)% ORAM READ

100% ORAM WRITE

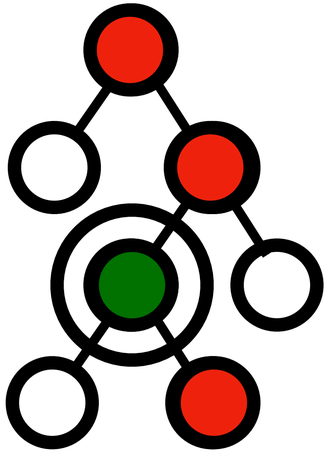

K% One-sided READ

Opportunity of One-sided Communication

OR

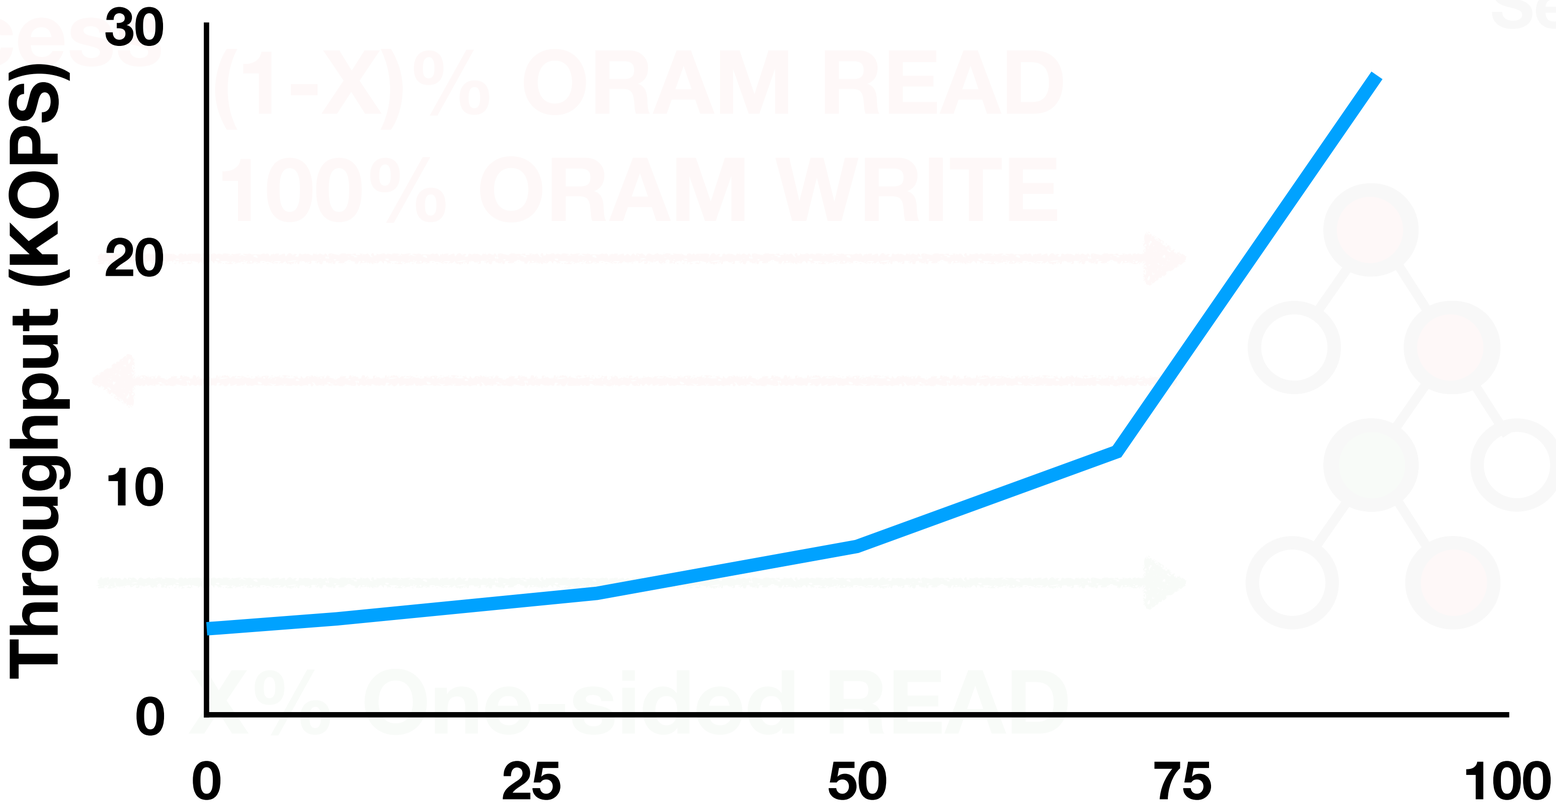

Conclusion

- **Security concerns** of one-sided communication
- Tradeoffs between **Performance** and **Security**
- **Hardware Vendor**, **Software Developers**, and **Datacenter**

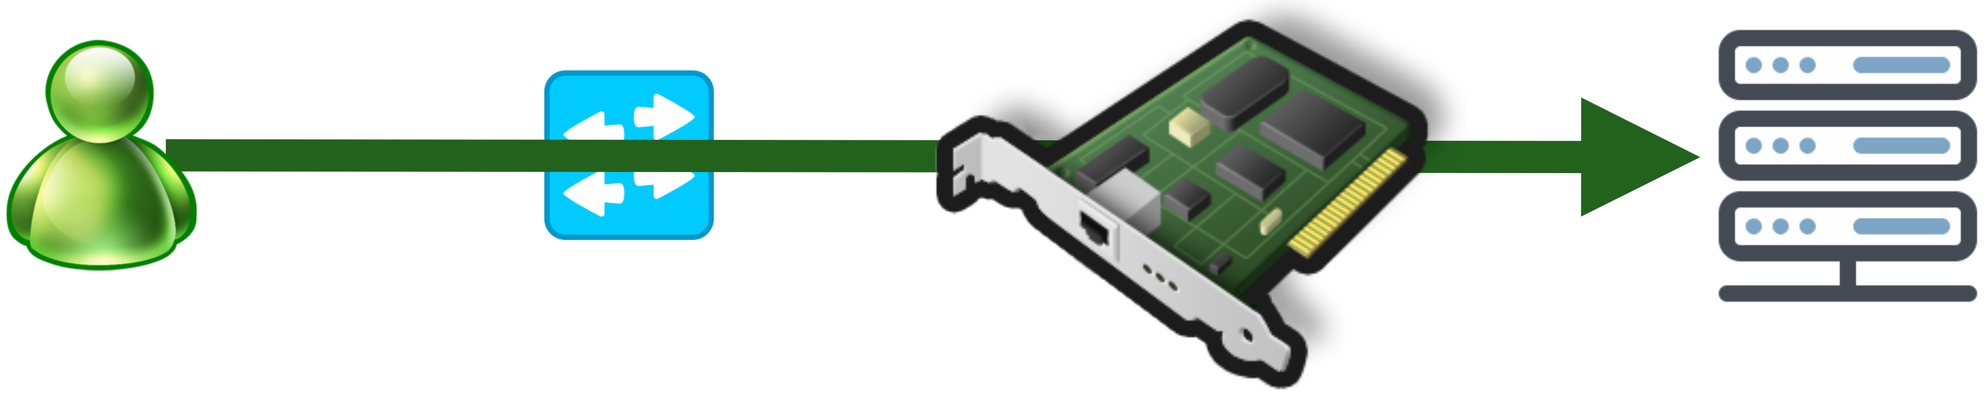

Conclusion

- **Security concerns** of one-sided communication
- Tradeoffs between **Performance** and **Security**
- **Hardware Vendor**, **Software Developers**, and **Datacenter**

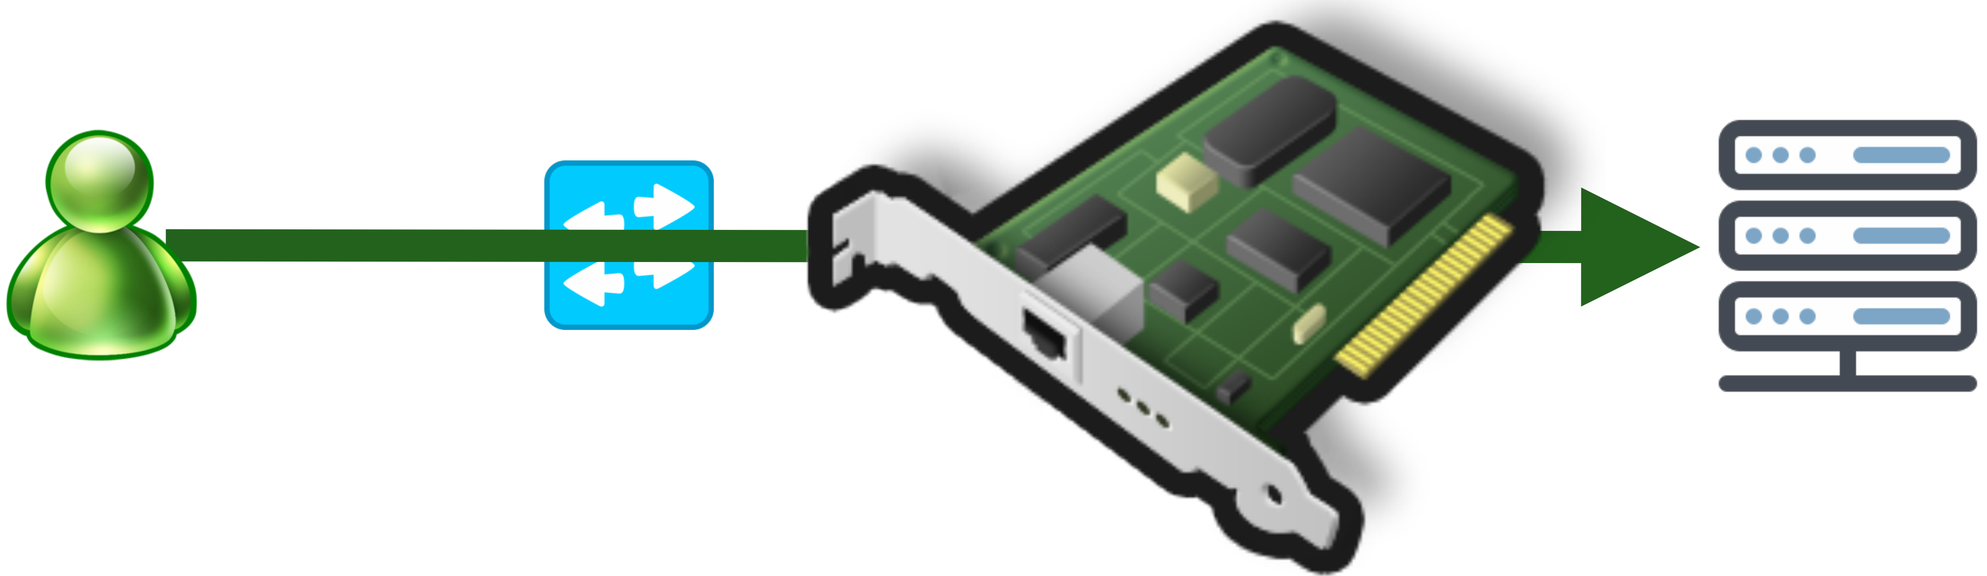

Thank you
Questions?

WukLab

wuklab.io

PURDUE UNIVERSITY