

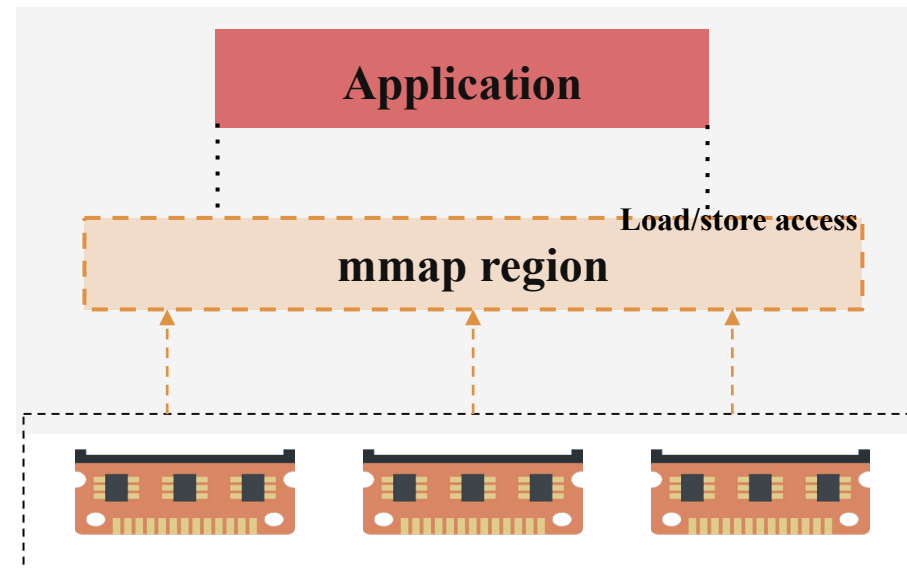
# TENET: Memory Safe and Fault Tolerant Persistent Transactional Memory

R. Madhava Krishnan, Diyu Zhou, **Wook-Hee Kim**, Sudarsun  
Kannan, Sanidhya Kashyap, Changwoo Min



# Boon and bane of Non-volatile Memory (NVM)

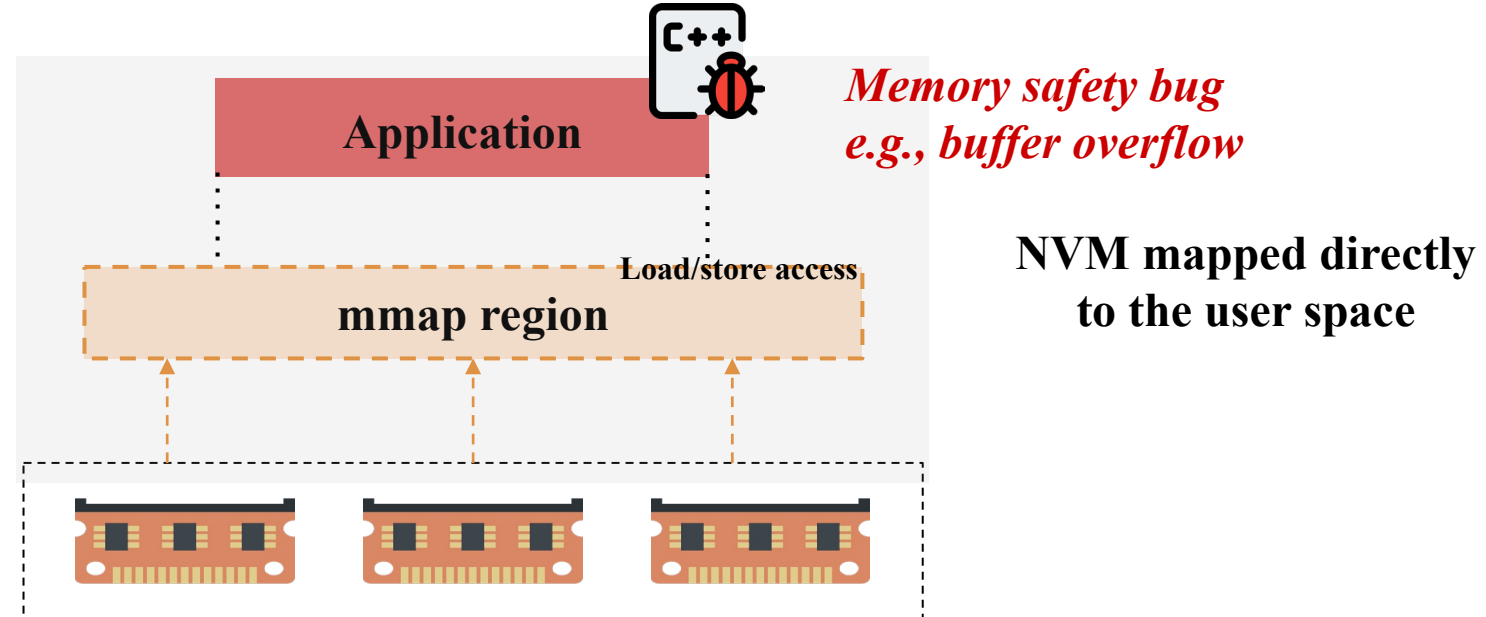
- Byte-addressability enables application to directly access NVM using load/store instructions
  - NVM is directly mapped to the application's address space



**NVM mapped directly  
to the user space**

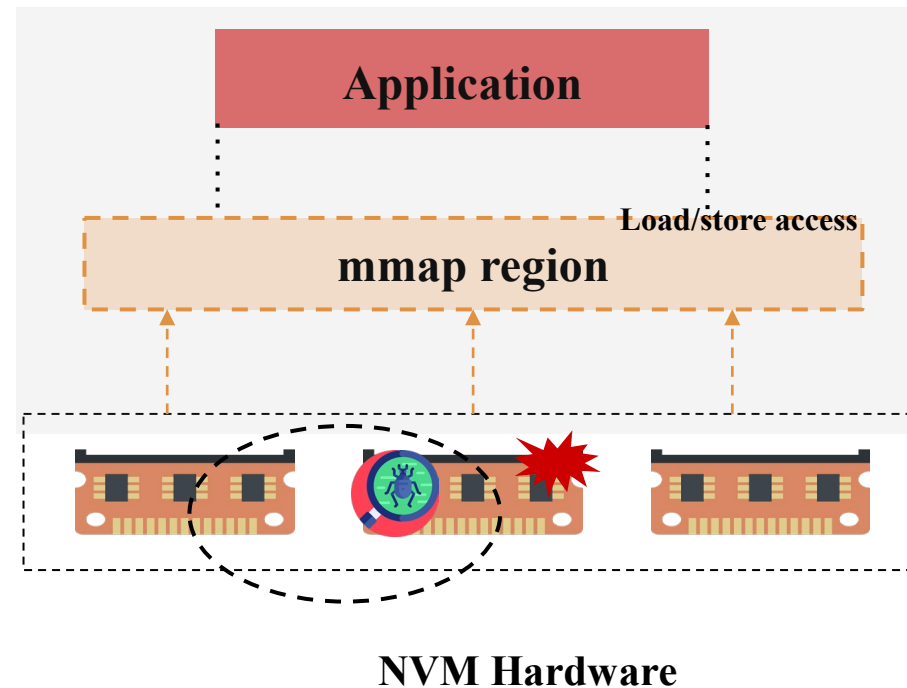
# Boon and bane of Non-volatile Memory (NVM)

- Byte-addressability makes NVM data vulnerable to memory safety bugs in the application



# Hardware (Media) errors are a threat too!

- **NVM data is vulnerable to Media Errors**
  - Device wear-out, power spikes, soft media faults etc



**Media errors corrupts the NVM data  
and the entire NVM page (data) is lost**

# Research problem that we tackle..

- How to detect memory safety bugs in the application and prevent it from corrupting the NVM data?
- How to prevent data loss due to the NVM media errors?

**TENET**

- Background: NVM memory safety errors
- TENET Overview
- TENET Design
- Evaluation
- Conclusion

# Types of memory safety violations

- Memory Safety Violations
  - Spatial Safety Violations
  - Temporal Safety Violations

## Spatial Safety Violations

```
memcpy(buff, src, 64)
```

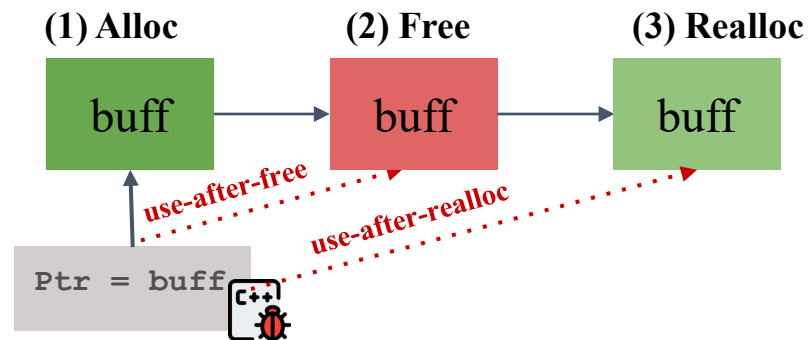


**buffer overflow**



Spatial safety violations happens when applications access the memory ***beyond the allocated range***

## Temporal Safety Violations



Temporal safety violations happens when applications access the memory ***using dangling pointers***

# Types of Media Errors

- NVM Media Errors
  - Correctable Media Errors
  - Uncorrectable Media Errors
- NVM has high Random Bit Error Rate (RBER)  $\approx$  NAND flash
- **Uncorrectable media errors (UME) are detected by the hardware ECC but can not be corrected**
  - UME can happen at random offset and the OS kernel offlines the corrupted NVM page
  - **Application is responsible for fixing the corrupted NVM page**

Applications are required to maintain a backup of NVM data to rollback the affected NVM page to prevent data loss



# Summary of prior Persistent Transactional Memory (PTM) works

PTM	Baseline PTM*	Spatial Safety	Temporal Safety	Fault Tolerance	Performance Overhead	NVM Cost Overhead
Libpmemobj- <b>R</b>	libpmemobj				<b>100%</b>	<b>High</b>
SafePM [Eurosys-22]	libpmemobj				<b>55%</b>	
Pangolin [ATC-19]	libpmemobj				<b>67%</b>	<b>Moderate</b>

Guaranteeing memory safety and fault tolerance at **a lower performance overhead** and **cost** is a very challenging problem

- Background: NVM memory safety errors
- TENET Overview
- TENET Design
- Evaluation
- Conclusion

# TENET overview: Goals and Assumptions

- Protect NVM data from a buggy application code
  - Guarantee spatial safety and temporal safety
- Protect NVM data against Uncorrectable Media Errors (UME)
  - Guarantee a performance and cost efficient fault tolerance
- Adversarial attacks are out-of-scope
- TENET library code and OS kernel are trusted (TCB)

TENET is a NVM programming framework to develop memory safe and fault tolerant NVM data structures and applications

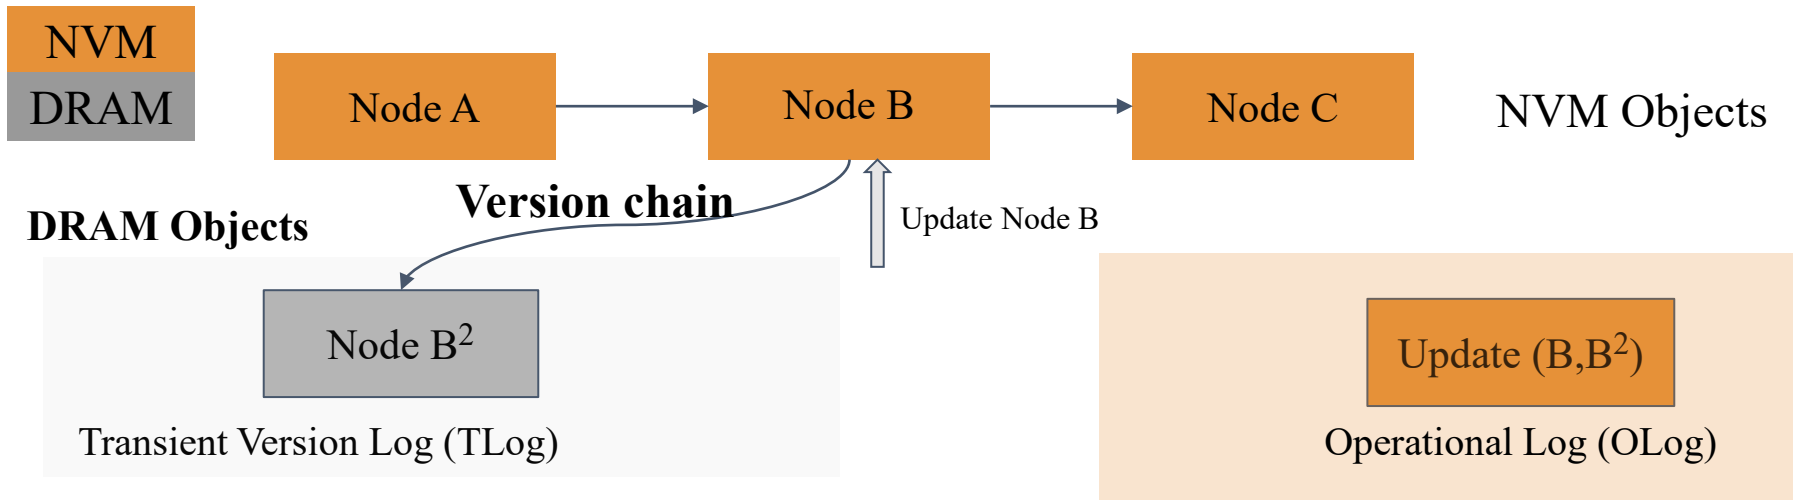
# TENET overview: Programming Model

- TENET provides persistent transaction programming model
  - TENET uses TimeStone<sup>[1]</sup> persistent transactional memory (PTM)
  - TimeStone is the state-of-the-art high-performing, highly scalable PTM
  - TimeStone does not provide memory safety or fault tolerance

[1] Durable Transactional Memory Can Scale with TimeStone, ASPLOS'20

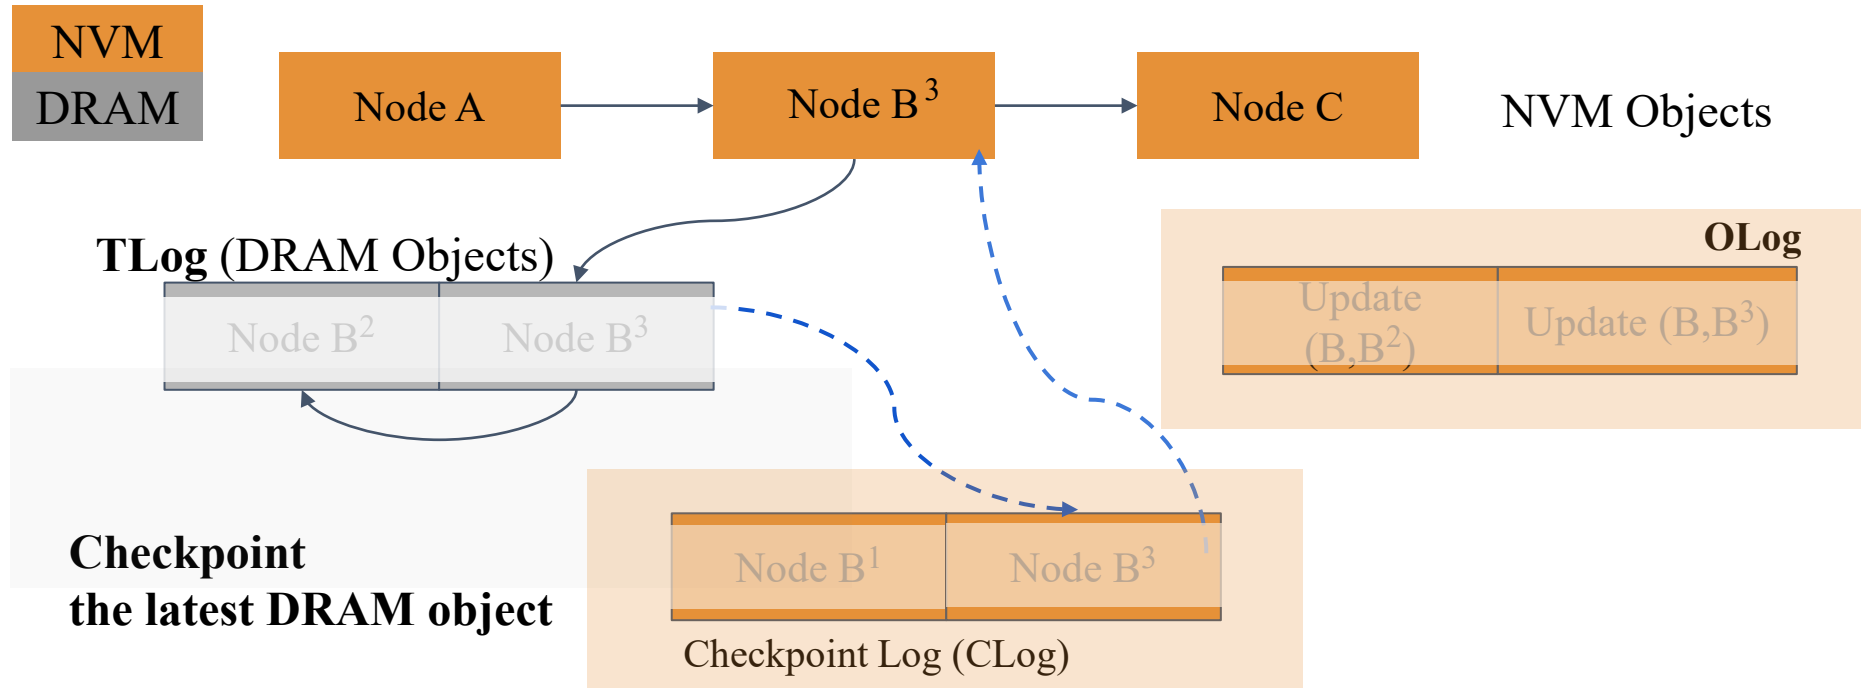
# Overview of TimeStone PTM

- TimeStone maintains the version chain of an NVM object on the DRAM



- Different versions of a NVM object are created in the TLog
- Operational logging to guarantee durability for the updates on TLog
- OLog will be replayed during the recovery to get back the DRAM objects

# Overview of TimeStone PTM



# Talk Outline

---

- Background: NVM memory safety errors
- TENET Overview
- TENET Design
- Evaluation
- Conclusion

# Spatial safety design in TENET

---

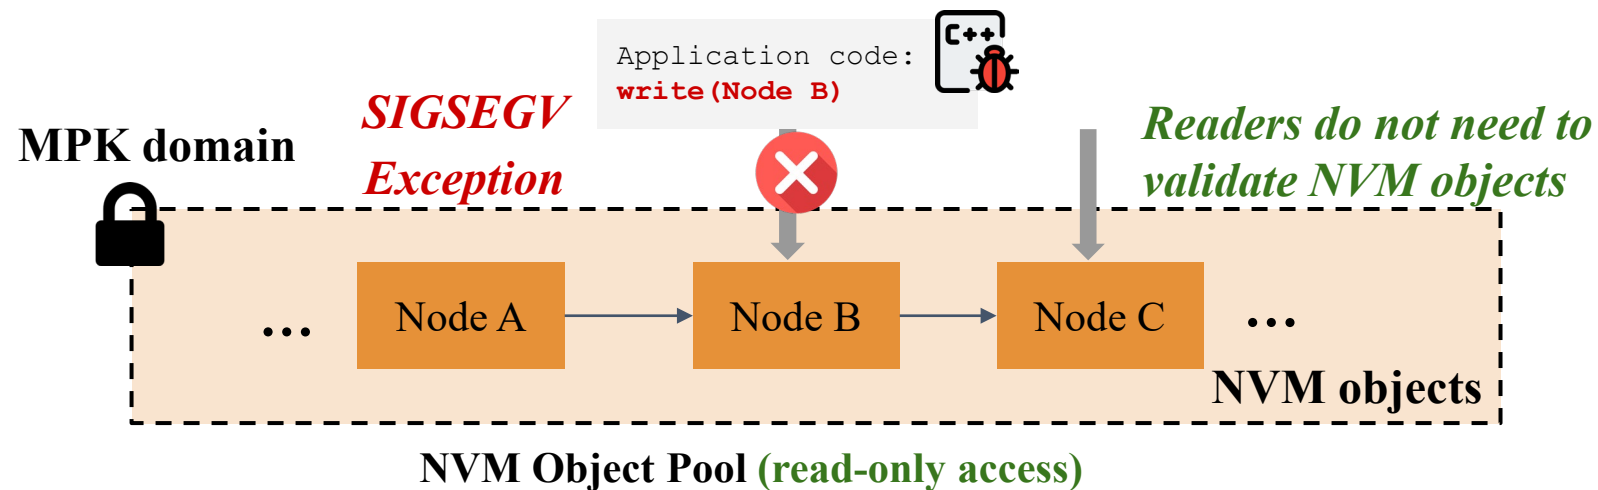
- Application code or any code outside the TENET library is **not allowed** to perform direct NVM writes
- Only the TENET library code **is allowed** to perform writes to the NVM data





# Prevent direct NVM writes using MPK

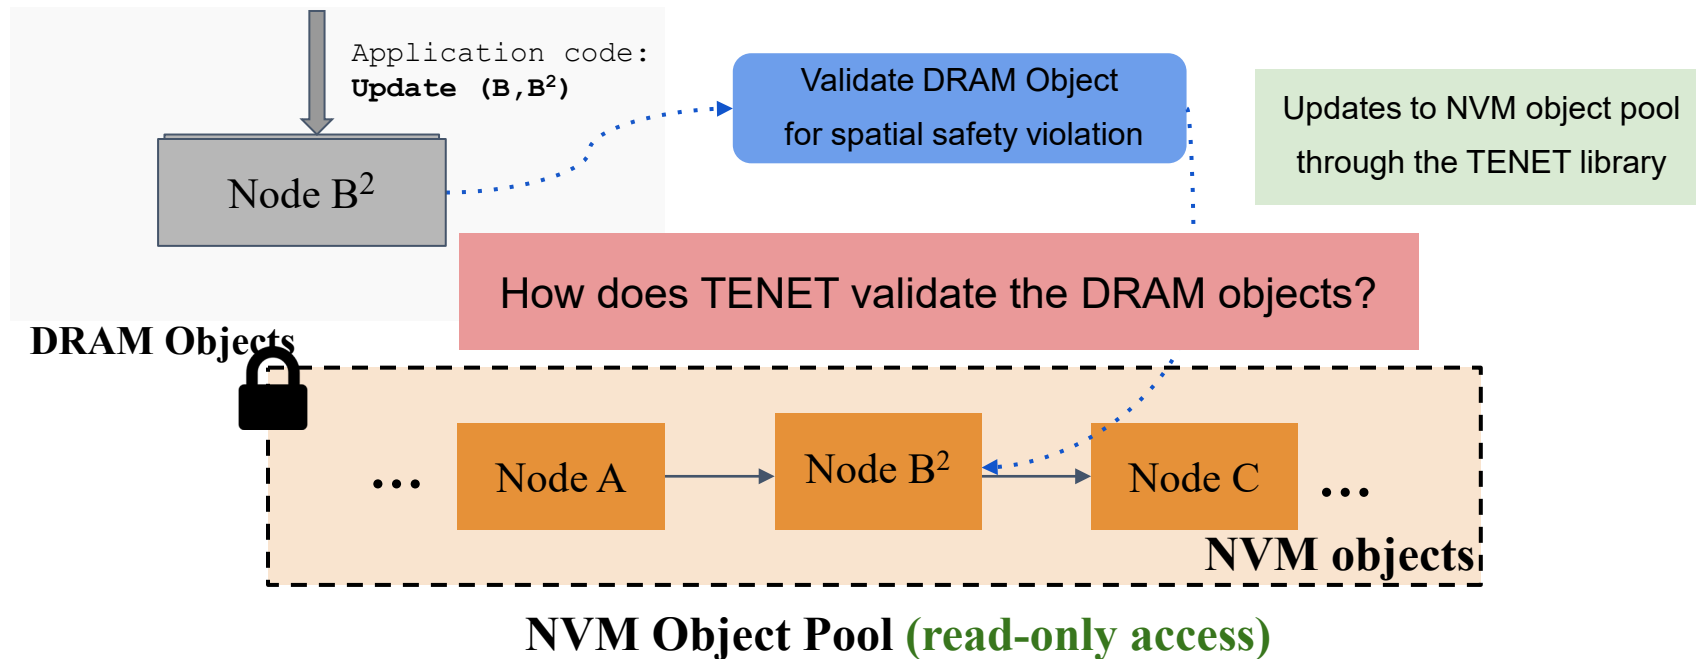
- TENET uses MPK to enforce read-only access to the NVM object pool for all the code outside of the TENET library



How does application writes to the NVM objects?

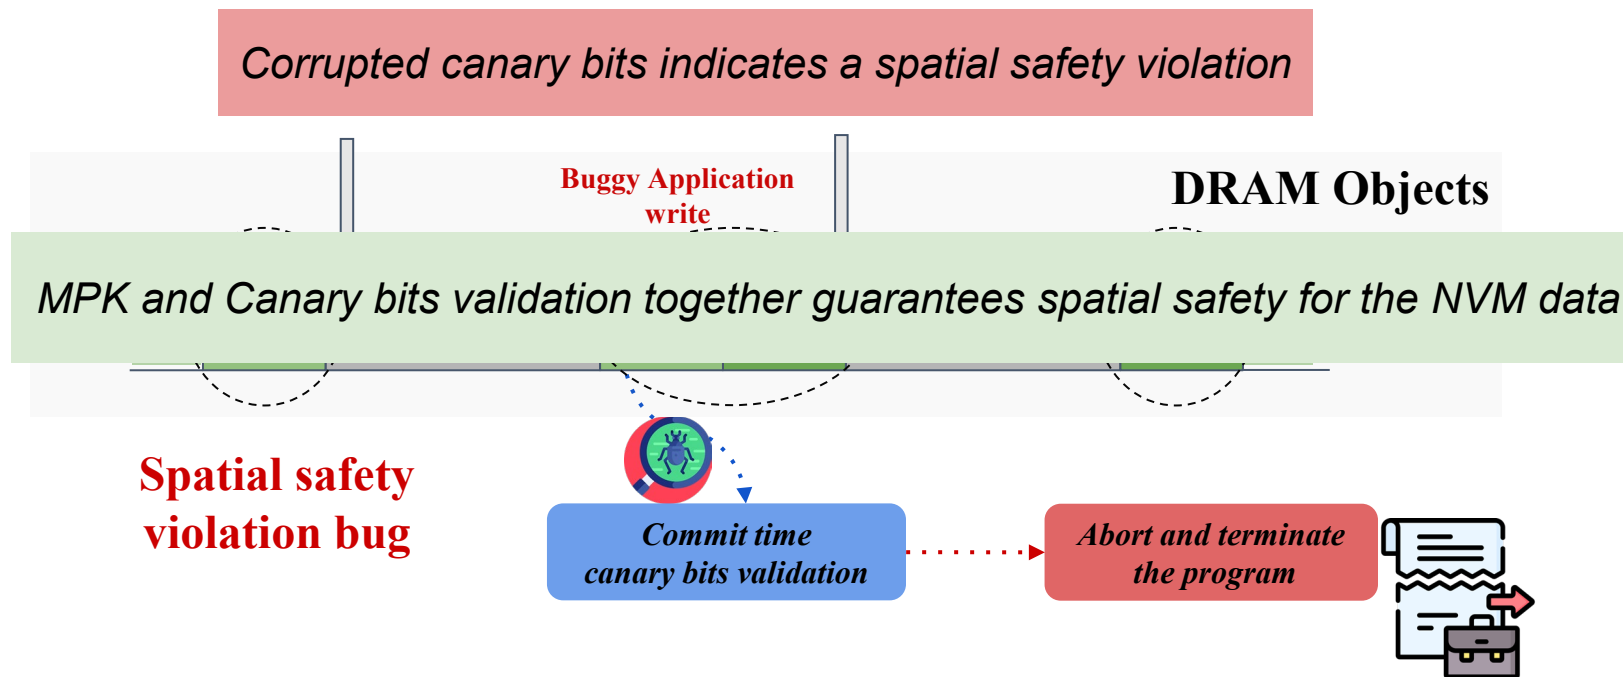
# Prevent direct NVM writes using MPK

- Application writes only on the DRAM region and TENET writes back the DRAM object to the NVM after validating it for spatial safety



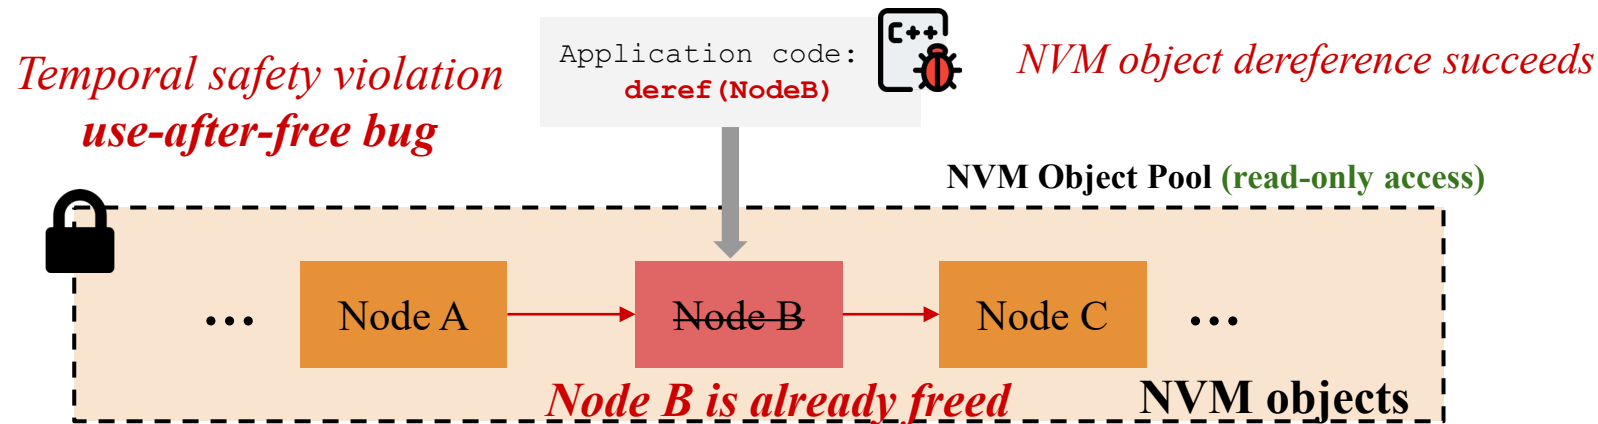
# Protecting DRAM objects using canary bits

- TENET assigns 8 byte canaries at the boundary of a DRAM object at the time of its creation
- Canary bits are inspected when the application commits its transaction



# Read-only NVM access can cause temporal safety violations

- Does making NVM read-only solve all the problems and prevent NVM data corruption?



*How does TENET enforce temporal memory safety for the NVM objects?*

# Enforcing temporal safety for NVM objects using pointer tags

- NVM address is tagged at the time of creation; the tag is stored in the allocated NVM object and a copy of the tag is encoded in the upper 16 bits of the NVM pointer



- Node B's address → 0x00001265FFCAB734; Tag → 0xCAFE
  - Encoded pointer → Node B || Tag << 48 → 0xCAFE1265FFCAB734
- Upper 16 bits are unused*
- tag bits | NVM address

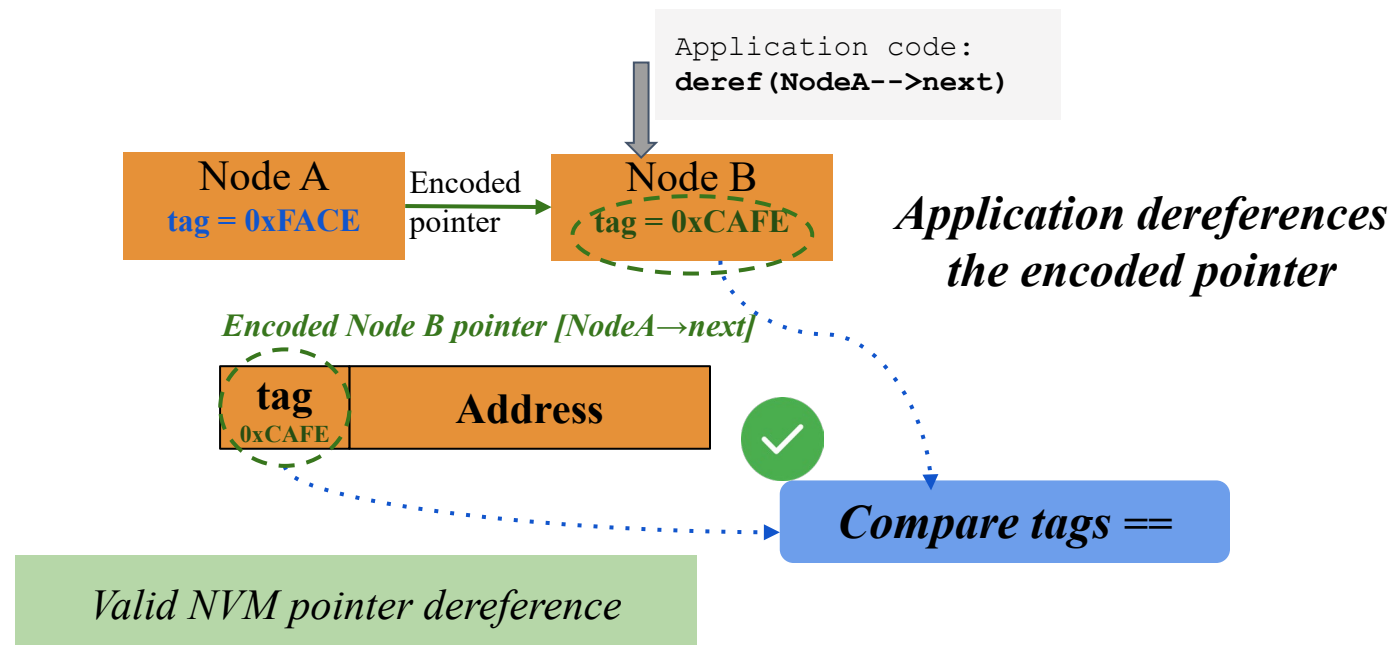
## Encoded pointer layout



*A copy of the tag is encoded to the upper 16 bits of Node B's address*

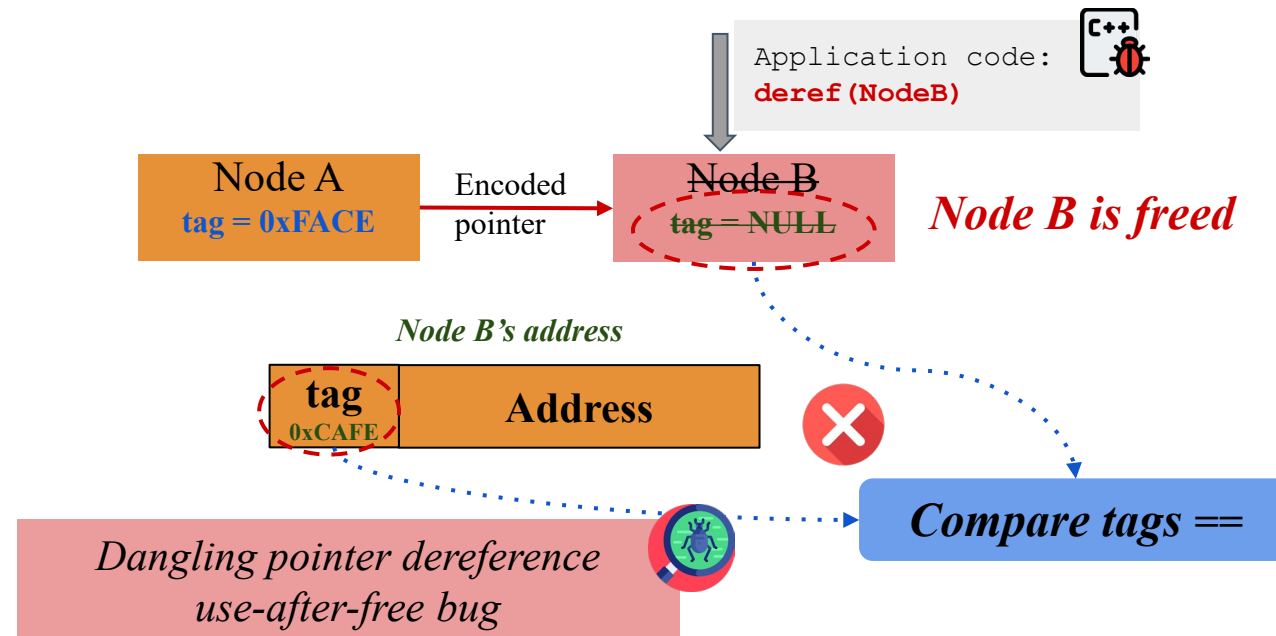
# Enforcing temporal safety for NVM objects using pointer tags

- Application accesses the NVM objects using the encoded pointer -- the encoded tag in the pointer is compared with the tag stored in the corresponding NVM object



# Enforcing temporal safety for NVM objects using pointer tags

- Dangling pointer is detected by comparing the tag stored in the NVM object with the tag encoded in the pointer to the NVM object





# Replicating NVM data for fault tolerance against UME

- NVM data corruption due to software errors
  - Spatial memory safety → MPK + canary bits validation ✓
  - Temporal memory safety → Pointer tags validation ✓

*How does TENET make the NVM data fault tolerance against the UME?*

# Replicating NVM data for fault tolerance against UME

- TENET replicates the NVM data to the local SSD to maintain backup copy
- Restore the corrupted NVM page from the SSD replica
- TENET's replication provides many desirable properties
  - **Cost efficiency** → replicating to the local SSD
  - **Performance efficiency** → replicating the data out-of-the critical path
  - **Consistent loss-less recovery**

Refer to the paper for more details

# Talk Outline

---

- Background: NVM memory safety errors
- TENET Overview
- TENET Design
- Evaluation
- Conclusion

- Evaluation Questions
  - How does TENET compare against the prior PTM works in terms of features and performance overhead?
  - How much overhead does TENET incurs over its baseline PTM system TimeStone?
- Evaluation Settings
  - We use a 2 socket server with 64 core Intel Xeon Gold CPU
    - 64GB DRAM, 512GB NVM, 1TB SSD
  - We evaluate two different versions of TENET
    - TENET-MS → supports only memory safety
    - TENET → supports memory safety and fault tolerance
  - We evaluate TENET with different data structures for different read/write ratios
    - YCSB workloads and microbenchmarks

# Comparison of TENET with the other PTMs

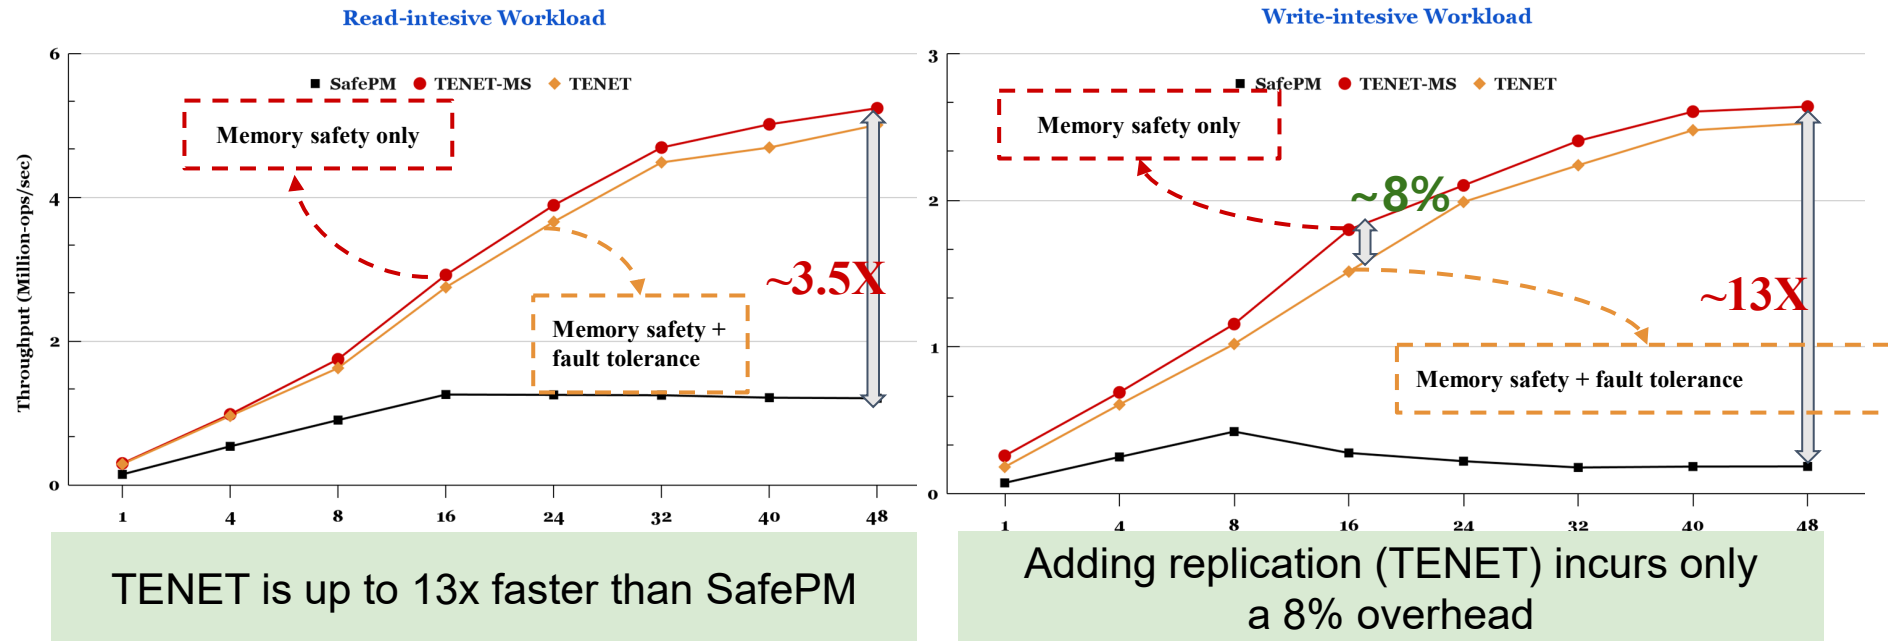
	PTM	Baseline PTM*	Spatial safety	Temporal Safety	Fault tolerance	
→	Libpmemobj-R	libpmemobj	✗	✗	✓	Replicates NVM data to a local NVM pool
→	SafePM [Eurosys-22]	libpmemobj	✓	✓	✗	Adds address sanitizer (Asan) to the libpmemobj
→	Pangolin [ATC-19]	libpmemobj	✓	✗	✓	Supports parity based replication and object checksums
→	TENET	TimeStone	✓	✓	✓	

TENET is the only PTM to provide spatial memory safety, temporal memory safety, and fault tolerance for the NVM data

\*PTM - persistent transactional memory


\*Libpmemobj is a transactional library in the PMDK

# Performance for a concurrent hash table



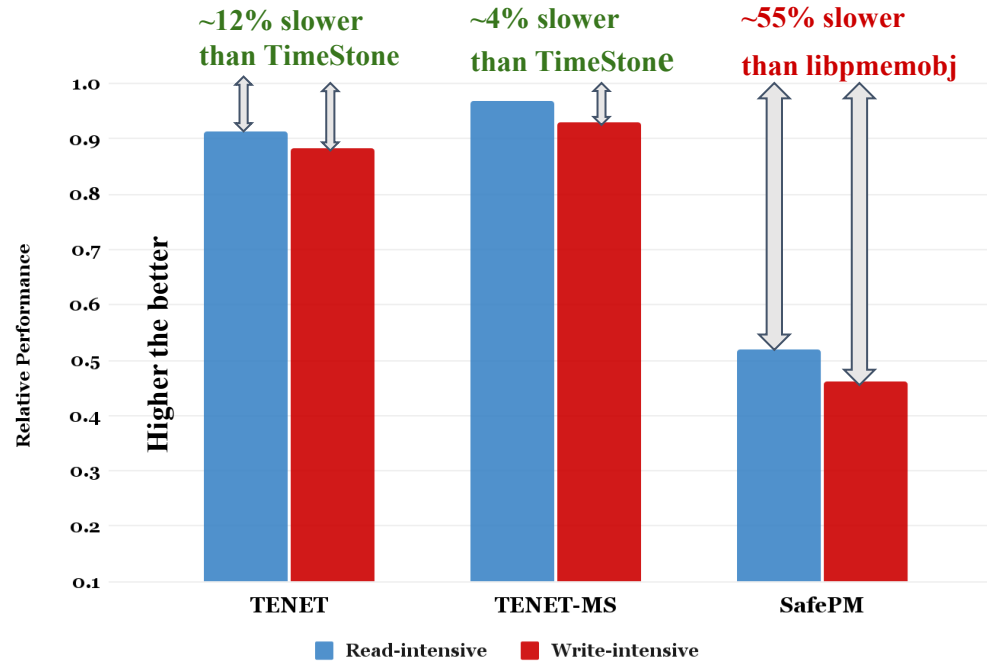
For a fair comparison let's compare the relative performance slowdown against their respective baseline PTM

# Performance for a concurrent hash table



- Performance is normalized to their respective baseline PTMs
  - SafePM normalized to the libpmemobj →  $\text{throughput (safePM)} / \text{throughput (libpmemobj)}$
  - TENET normalized to the TimeStone →  $\text{throughput (TENET)} / \text{throughput (TimeStone)}$

# Performance for a concurrent hash table



TENET does not require additional crash consistency operations for its memory safety metadata

- MPK → hardware primitive
- Pointer tags → embedded directly into the object

TENET does not perform memory safety validation for every NVM access

- Spatial safety checks performed only at the commit time
- Temporal safety checks performed only at the first-dereference of an NVM object

- Performance is normalized to their respective baseline PTMs
- SafePM normalized to the libpmemobj
- TENET normalized to the TimeStone

Refer to the paper for more details on these optimizations



- NVM is vulnerable to data corruption due to software bugs and media errors
- TENET is a NVM programming framework to design memory safe and fault tolerant NVM data structures and applications
  - **Spatial memory safety** → Memory protection keys (MPK) + Canary bits validation
  - **Temporal memory safety** → Encoded pointer tag validation during dereference
  - TENET guarantees **fault tolerance** for NVM data against uncorrectable media errors (UME)
    - Replicates the NVM objects to the local SSD
  - **TENET guarantees a robust memory protection and fault tolerance at a modest performance overhead**