
# Disk-Enabled Authenticated Encryption

Kevin Butler, Stephen McLaughlin, and Patrick McDaniel
Penn State University
MSST 2010, Incline Village, NV
6 May 2010

# Data Loss

- An increasingly large problem with disk storage
  - 64,000 records/SSNs of Ohio employees on media stolen from an intern's car (2007)
    - Governer: "What we're doing here is cautionary", no evidence of breach
    - Later findings: over 800,000 records stolen including those of regular citizens
  - 300,000 mental health histories on laptop stolen from PA public welfare department
  - 100,000 employee records on laptop lost by TSA
  - 3,400 National Guard employee records on stolen disk

# Integrity

- What happens when the media is retrieved?
  - What's been done to it?
- Confidentiality alone is not the answer
- Requirement: provide *integrity* as well as confidentiality for stored data
- Solution: authenticated encryption allows preservation of integrity and confidentiality
  - IEEE P1619.1

# Metadata Storage

- Regardless of mode of usage, requires MAC for integrity tag in addition to ciphertext storage
- Problem: ciphertext can be length-preserving, but integrity tags are not
  - Where to store additional data?
  - Not just MACs, but initialization vectors as well
  - 128-bit MAC, 96-bit IV

# On-Disk AE

- Proposal: store authentication material in NVRAM on the disk
  - Benefit: spatial locality of information and reduction of TCB compared to external metadata server
- What is the storage cost?
  - 1 TB disk and 512-byte sectors, = **54 GB** of NVRAM
  - Mitigate cost with *integrity sets* of adjacent sectors used for MAC calculation

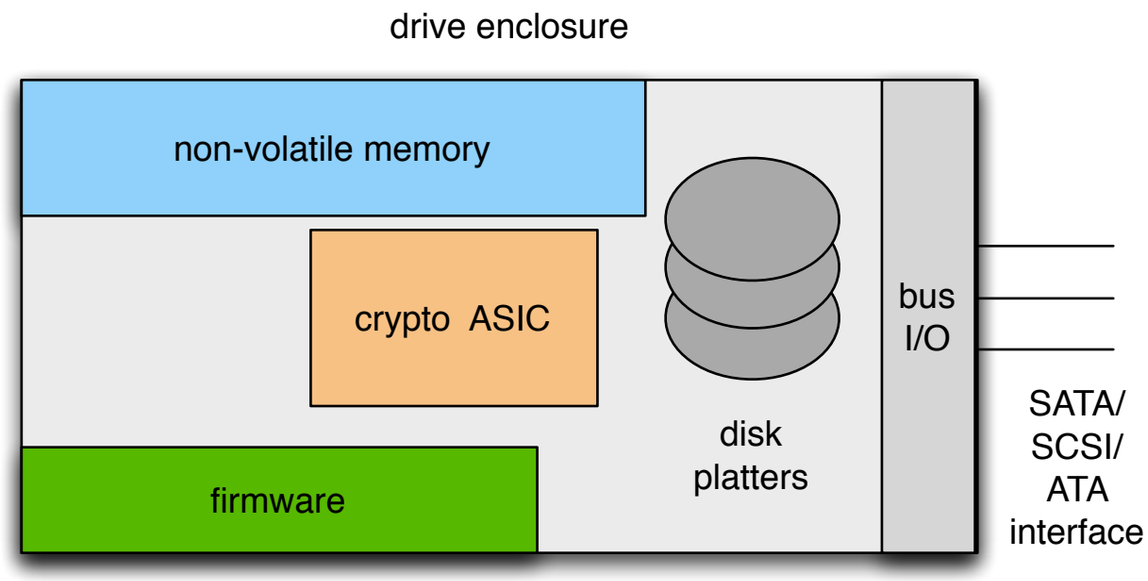

# Emulation of Disk AE

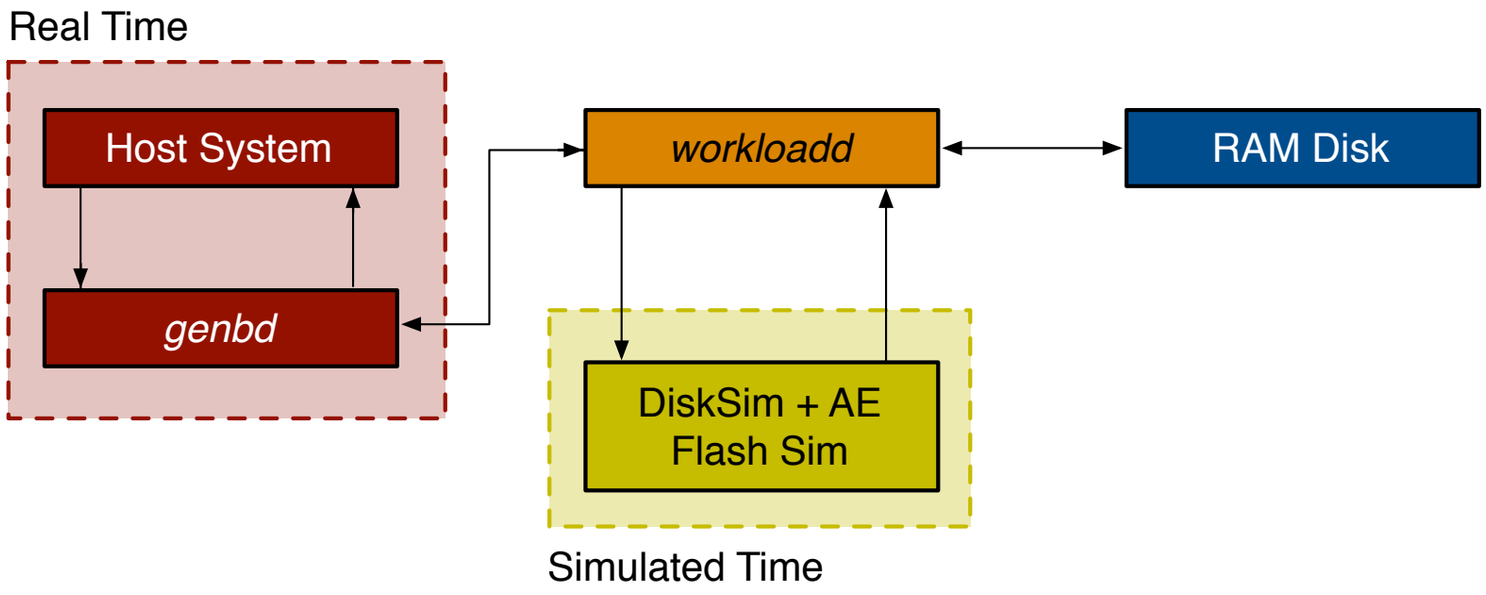


- Emulation vs. simulation: allows more accurate reflection of workloads since act as part of system while being easier than full implementation
- *workloadd* interfaces with Disksim in an event-timing loop (similar to the Memulator)
  - simulation events are handled faster but held back until they match wall-clock time to provide consistency

# Integrity Set Evaluation

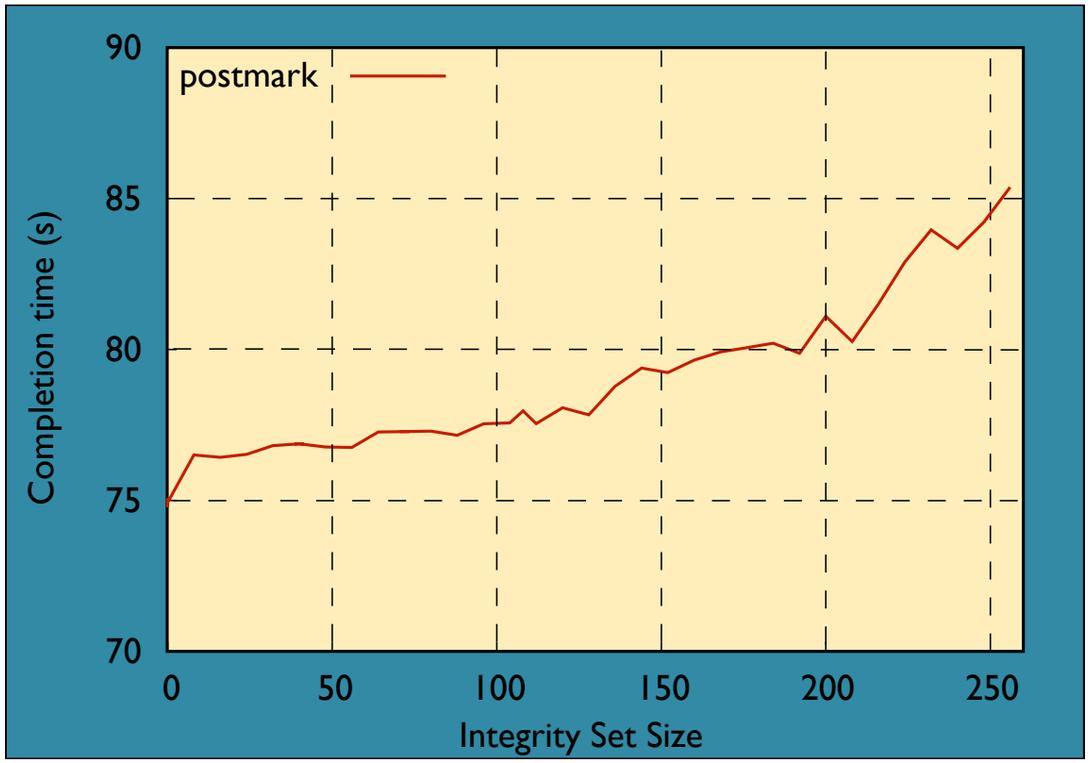


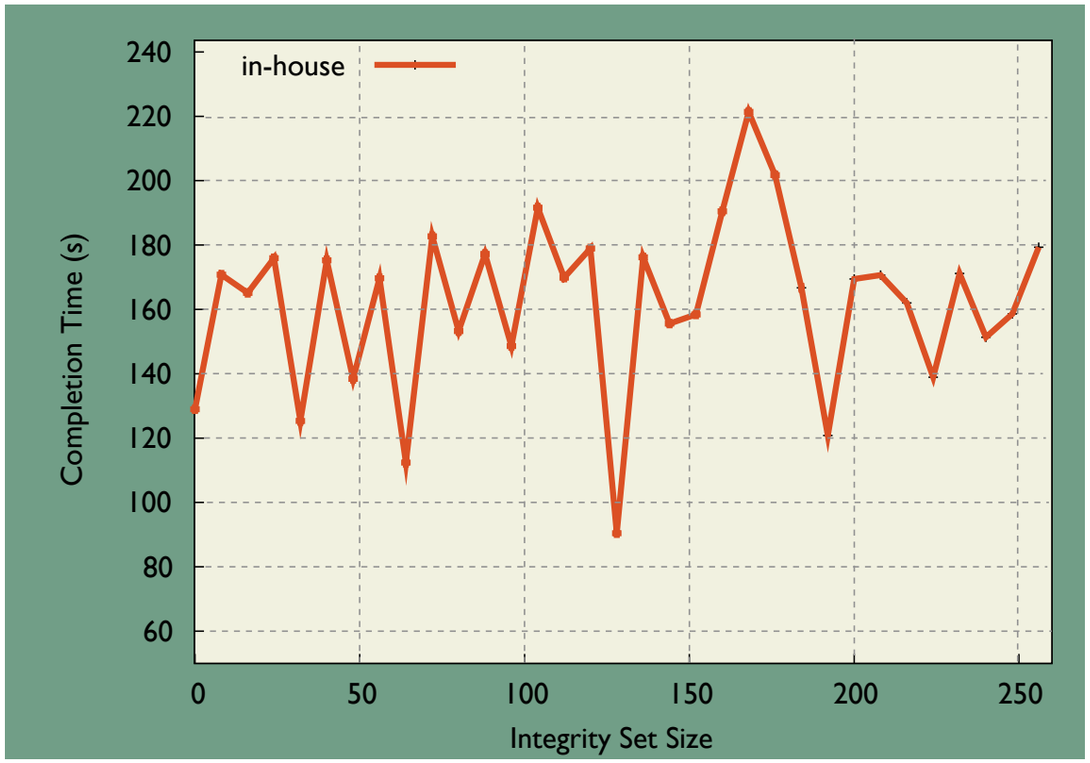


- Random workloads: increasing integrity set size increases completion time
  - rate not particularly high because transfer time does not appreciably increase
- Larger requests are influenced by track layout
- Also considered throughput (details available offline)

# Future Work

- Investigate new (more modern) DiskSim models
- Look at effects of on-disk metadata
- Understand effects of NVRAM metadata writes on overall reliability
- Investigate use in larger-scale storage systems

*Questions?* butler@cse.psu.edu