Autonomously Secure Disks

KEVIN BUTLER, STEPHEN MCLAUGHLIN, PATRICK MCDANIEL, AND YOUNGJAE KIM

Catastrophic Data Loss is Common

Between early August and early September 2007, there were 19 occurrences of data loss at organizations such as universities, hospitals, financial and health institutions, and government agencies including the military. All of these were caused by either improper access to data, improper disposal of storage, or by theft of devices.

State of Ohio: Intern took home a storage tape as part of backup protocol; tape stolen from car contained records of over 800,000 people on it.

Pennsylvania Public Welfare Department: Stolen laptop contained mental health histories of over 300,000 people.

Gap: Stolen laptop contained information on over 800,000 job applicants.

Multiboot Systems

Full disk encryption alone will not help systems that are multiboot capable. These systems do not enforce integrity, only confidentiality.

In addition, if multiple operating systems are running on a single disk, a successful attack against any operating system exposes the entire storage space, regardless of what operating system protections other OSes may have in place. Rootkits can also pose problems, particularly if they are VM-based. ***Whoever controls the machine's lowest layer will control the machine.***

New Hard Disk Architectures

We propose a new architecture for disk security, termed *autonomously-secure disks* (ASDs). These enforce a security perimeter at the external I/O interface, shrinking the effective TCB and reducing the potential attack surface for storage. As shown by the diagram, many components necessary for ASDs are available in current disk drives.

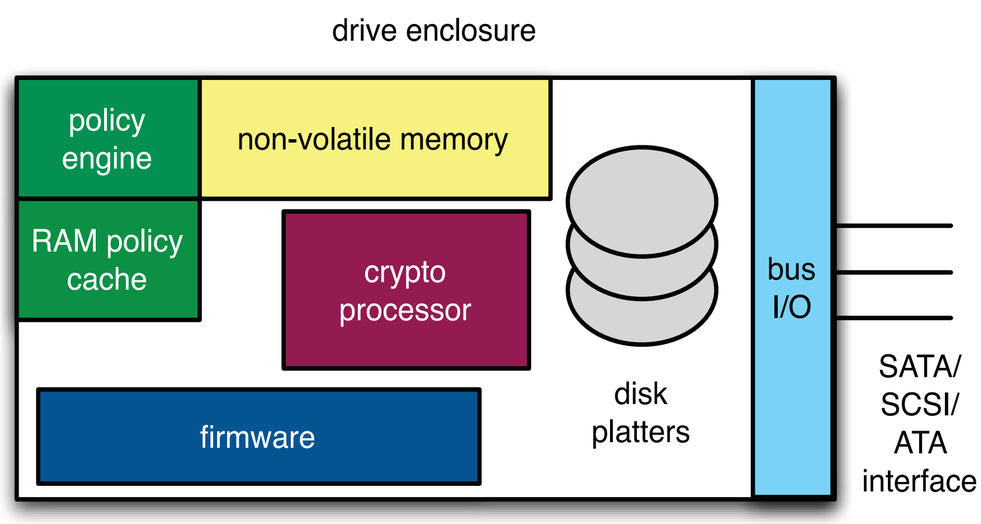

ASDs can support a variety of new applications that formerly were not possible within the disk itself.

Authenticated Encryption: In accordance with IEEE specification P1619.1, both confidentiality and integrity can be maintained on the drive, by leveraging the non-volatile memory on the disk to store HMACs and initialization vectors (IVs) on a block (or group of blocks) basis.

Capabilities: This arcahitecture allows disks to store capabilities in the non-volatile memory and access them on a block basis, while offloading capability management and delegation to the system.

Information Flow Preservation: ASDs can store data relating to security labels on drives in non-volatile memory and use on-board policy enforcement to make decisions about access rights at the block level, potentially supporting many information flow models such as Bell-La Padula and Clark-Wilson, and multiple policy architectures. Access decisions can be made faster through the addition of an on-board policy cache.

Block Integrity Sets

Storing HMAC and IV informtion in a drive's NVRAM is costly. It would require 54 GB of memory to store this information for a 1 TB drive. Additionally, we do not want to store this information on disk because of the need for fast, parallel access to metadata for policy decisions (e.g., access decisions for information flow labels). We thus require methods of mitigating NVRAM storage costs. We group multiple blocks of data together in an access unit that we call an *integrity set*.

We simulate the effect of integrity sets on our ASD architecture using a custom block driver, *genbd*, to move requests from kernel to user space, and a custom flash memory simulator. We drive workloads at the DiskSim simulator for emulating real-world drive mechanics.

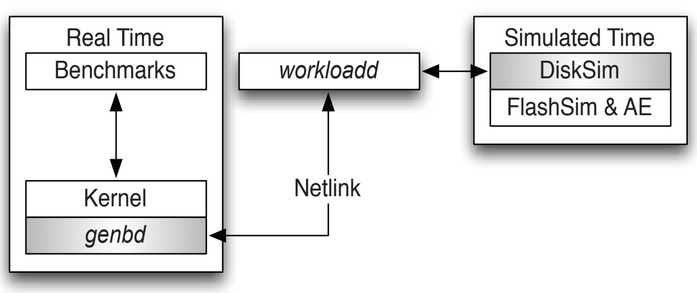

Evaluating Performance

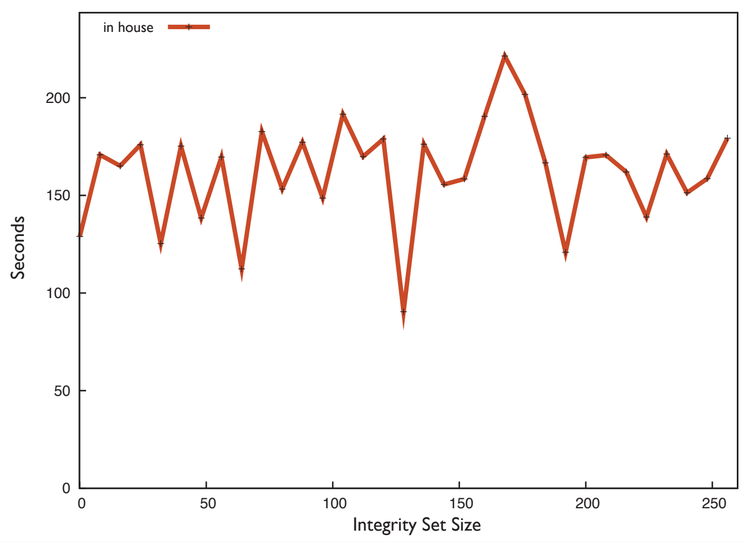

We evaluated our architecture with the PostMark benchmark (for random access) and a contiguous workload evaluator (for sequential access). While completion time increased as integrity set sizes increased, we found that set alignment played the predominant role in completion time for contiguous workloads, as shown on the left.

As shown on the right, we extrapolated our results by linear regression to a 1 TB drive and found that with an integrity set size of 16, only 4 GB of NVRAM is required while only affecting performance by about 2%.

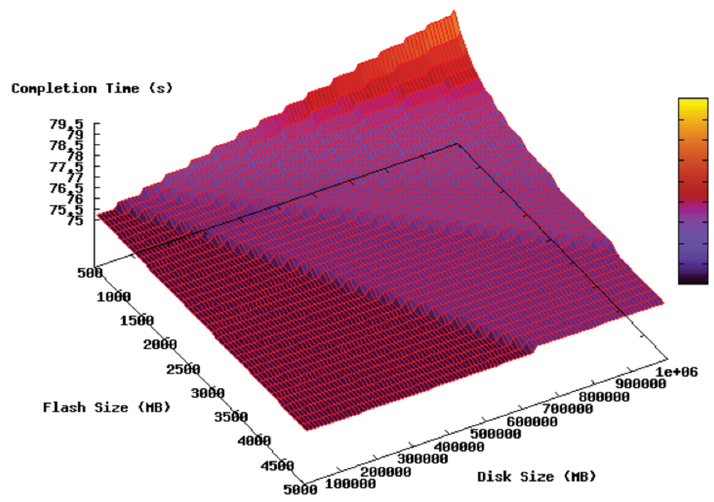

RESEARCH SUPPORTED BY NSF