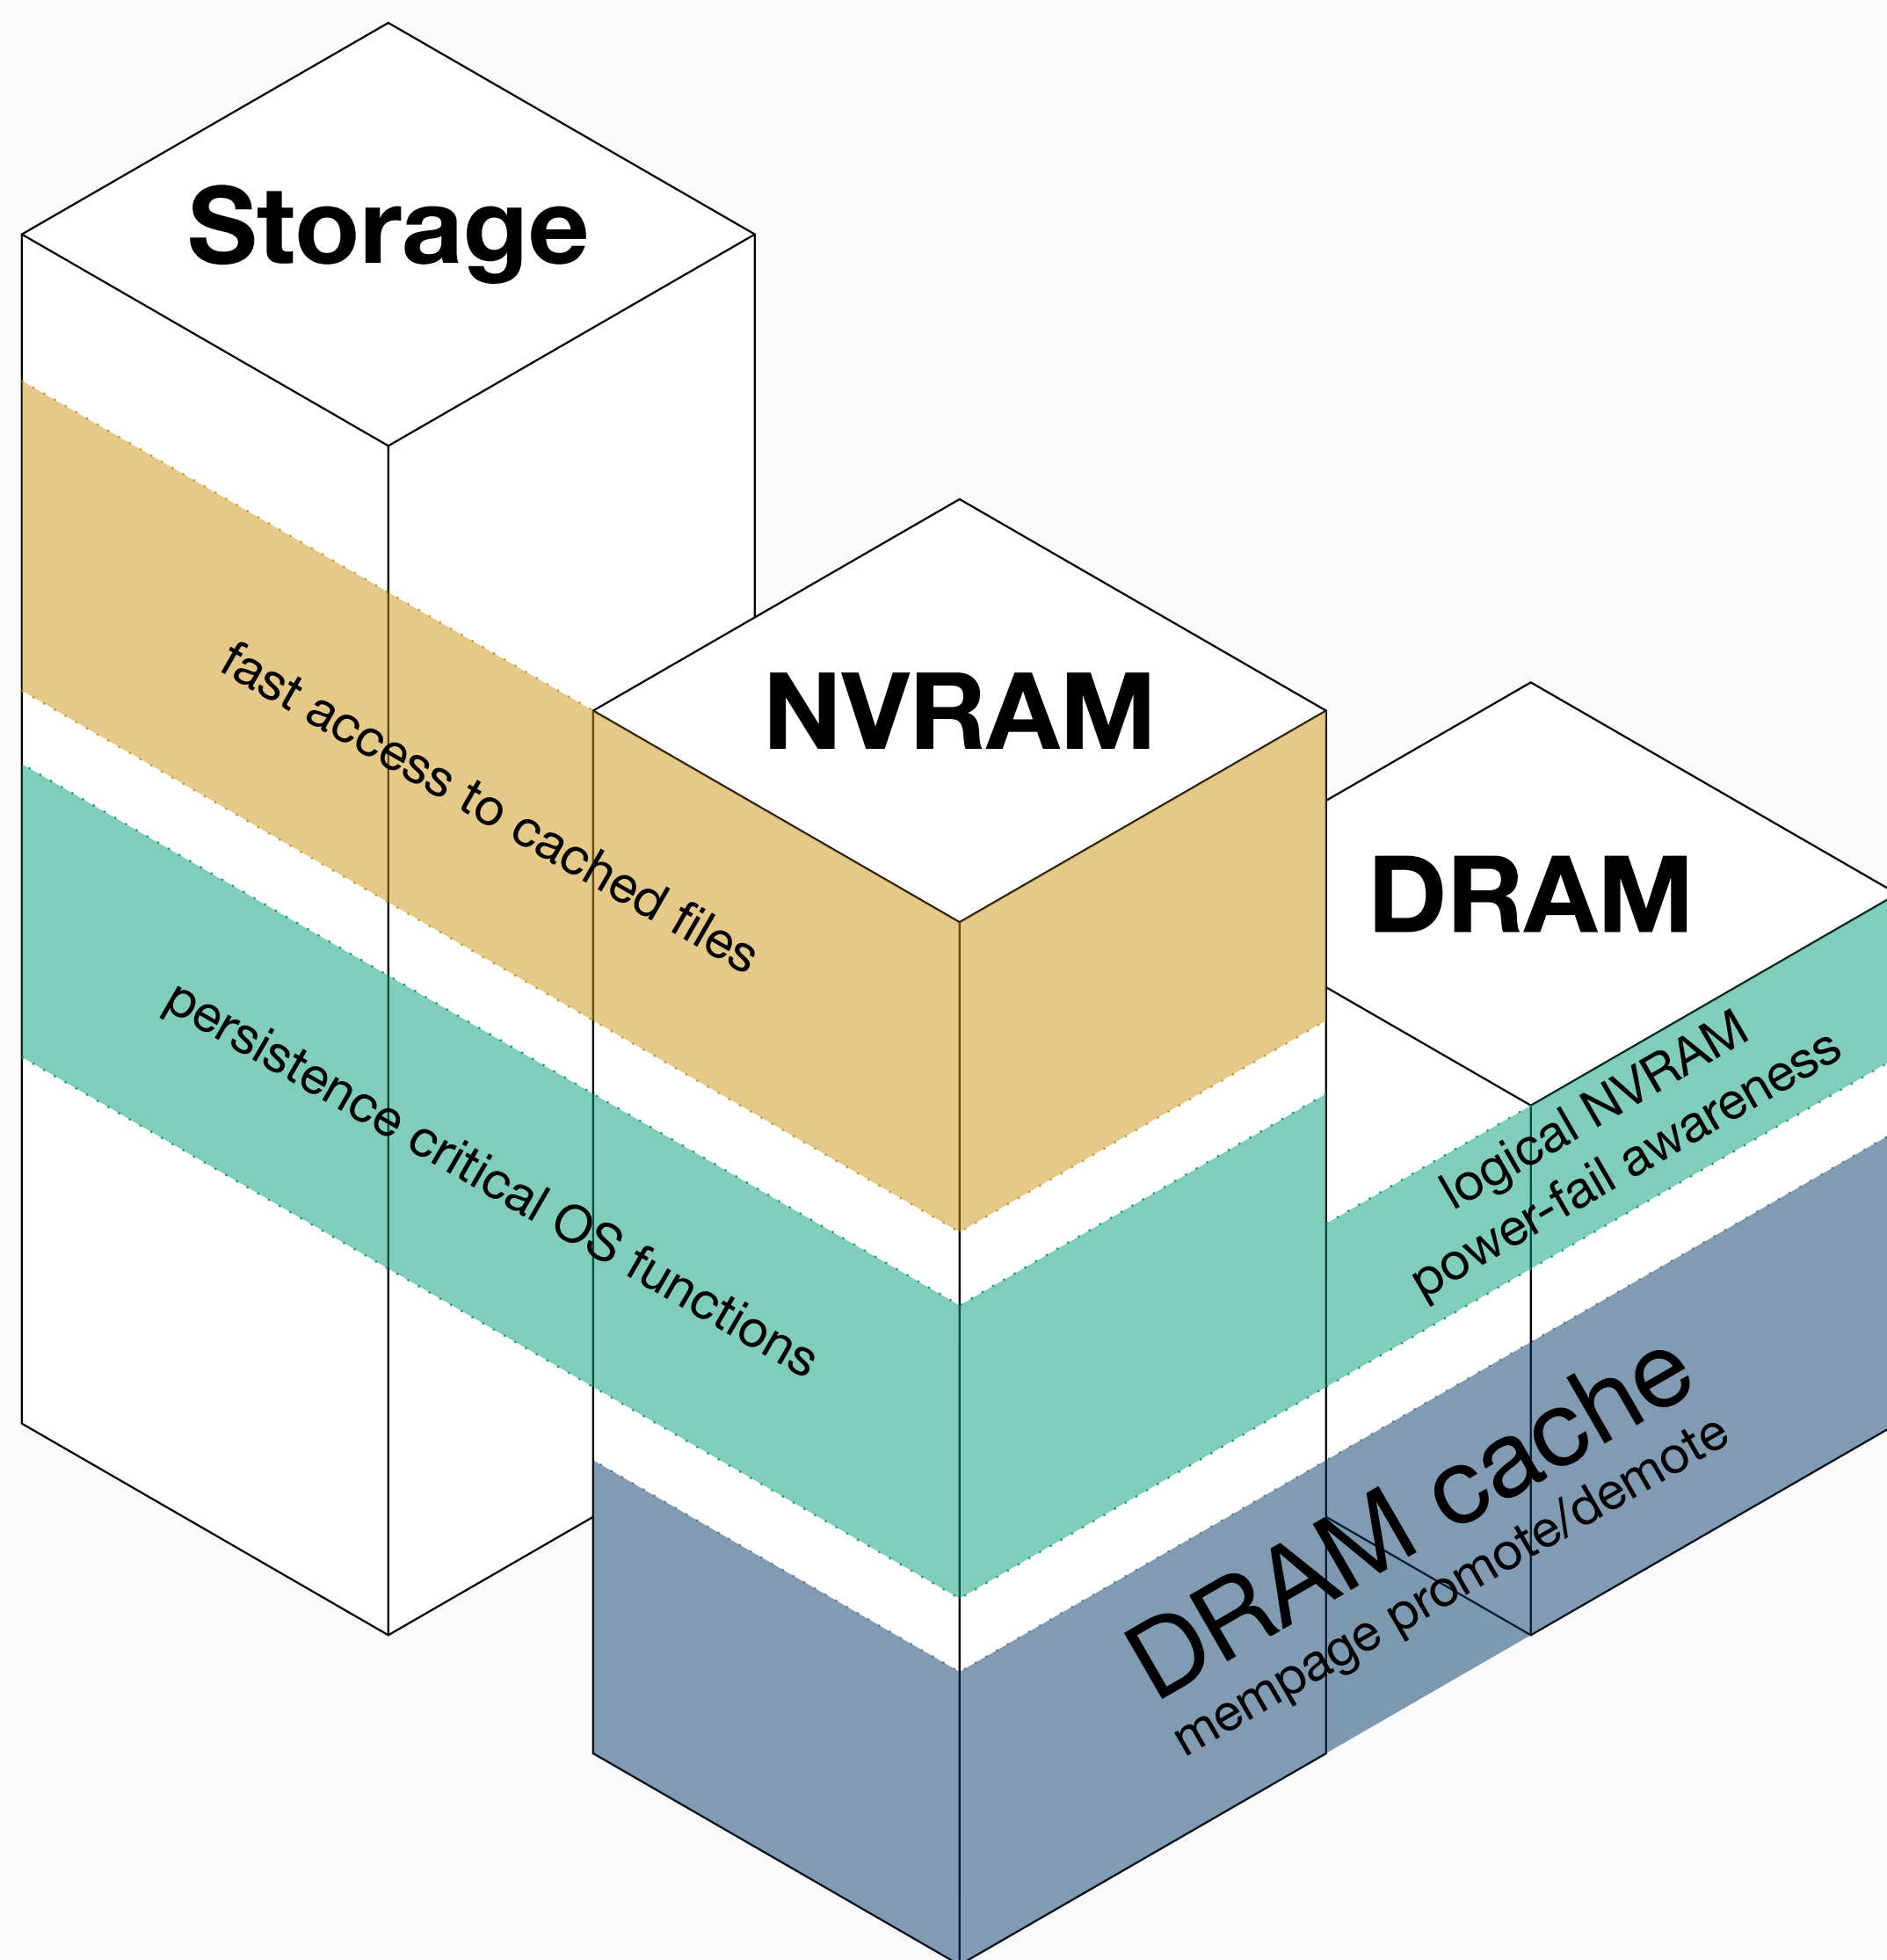


PAVE: Power-Fail Aware Byte-Addressable Virtual Non-Volatile Memory

Oliver Giersch, Dustin Nguyen, Jörg Nolte, Wolfgang Schröder-Preikschat

Memory Overview



Problem Statement

NVRAM is significantly faster than other types of storage, capacity is pretty high, and no energy is needed to maintain its state. But:

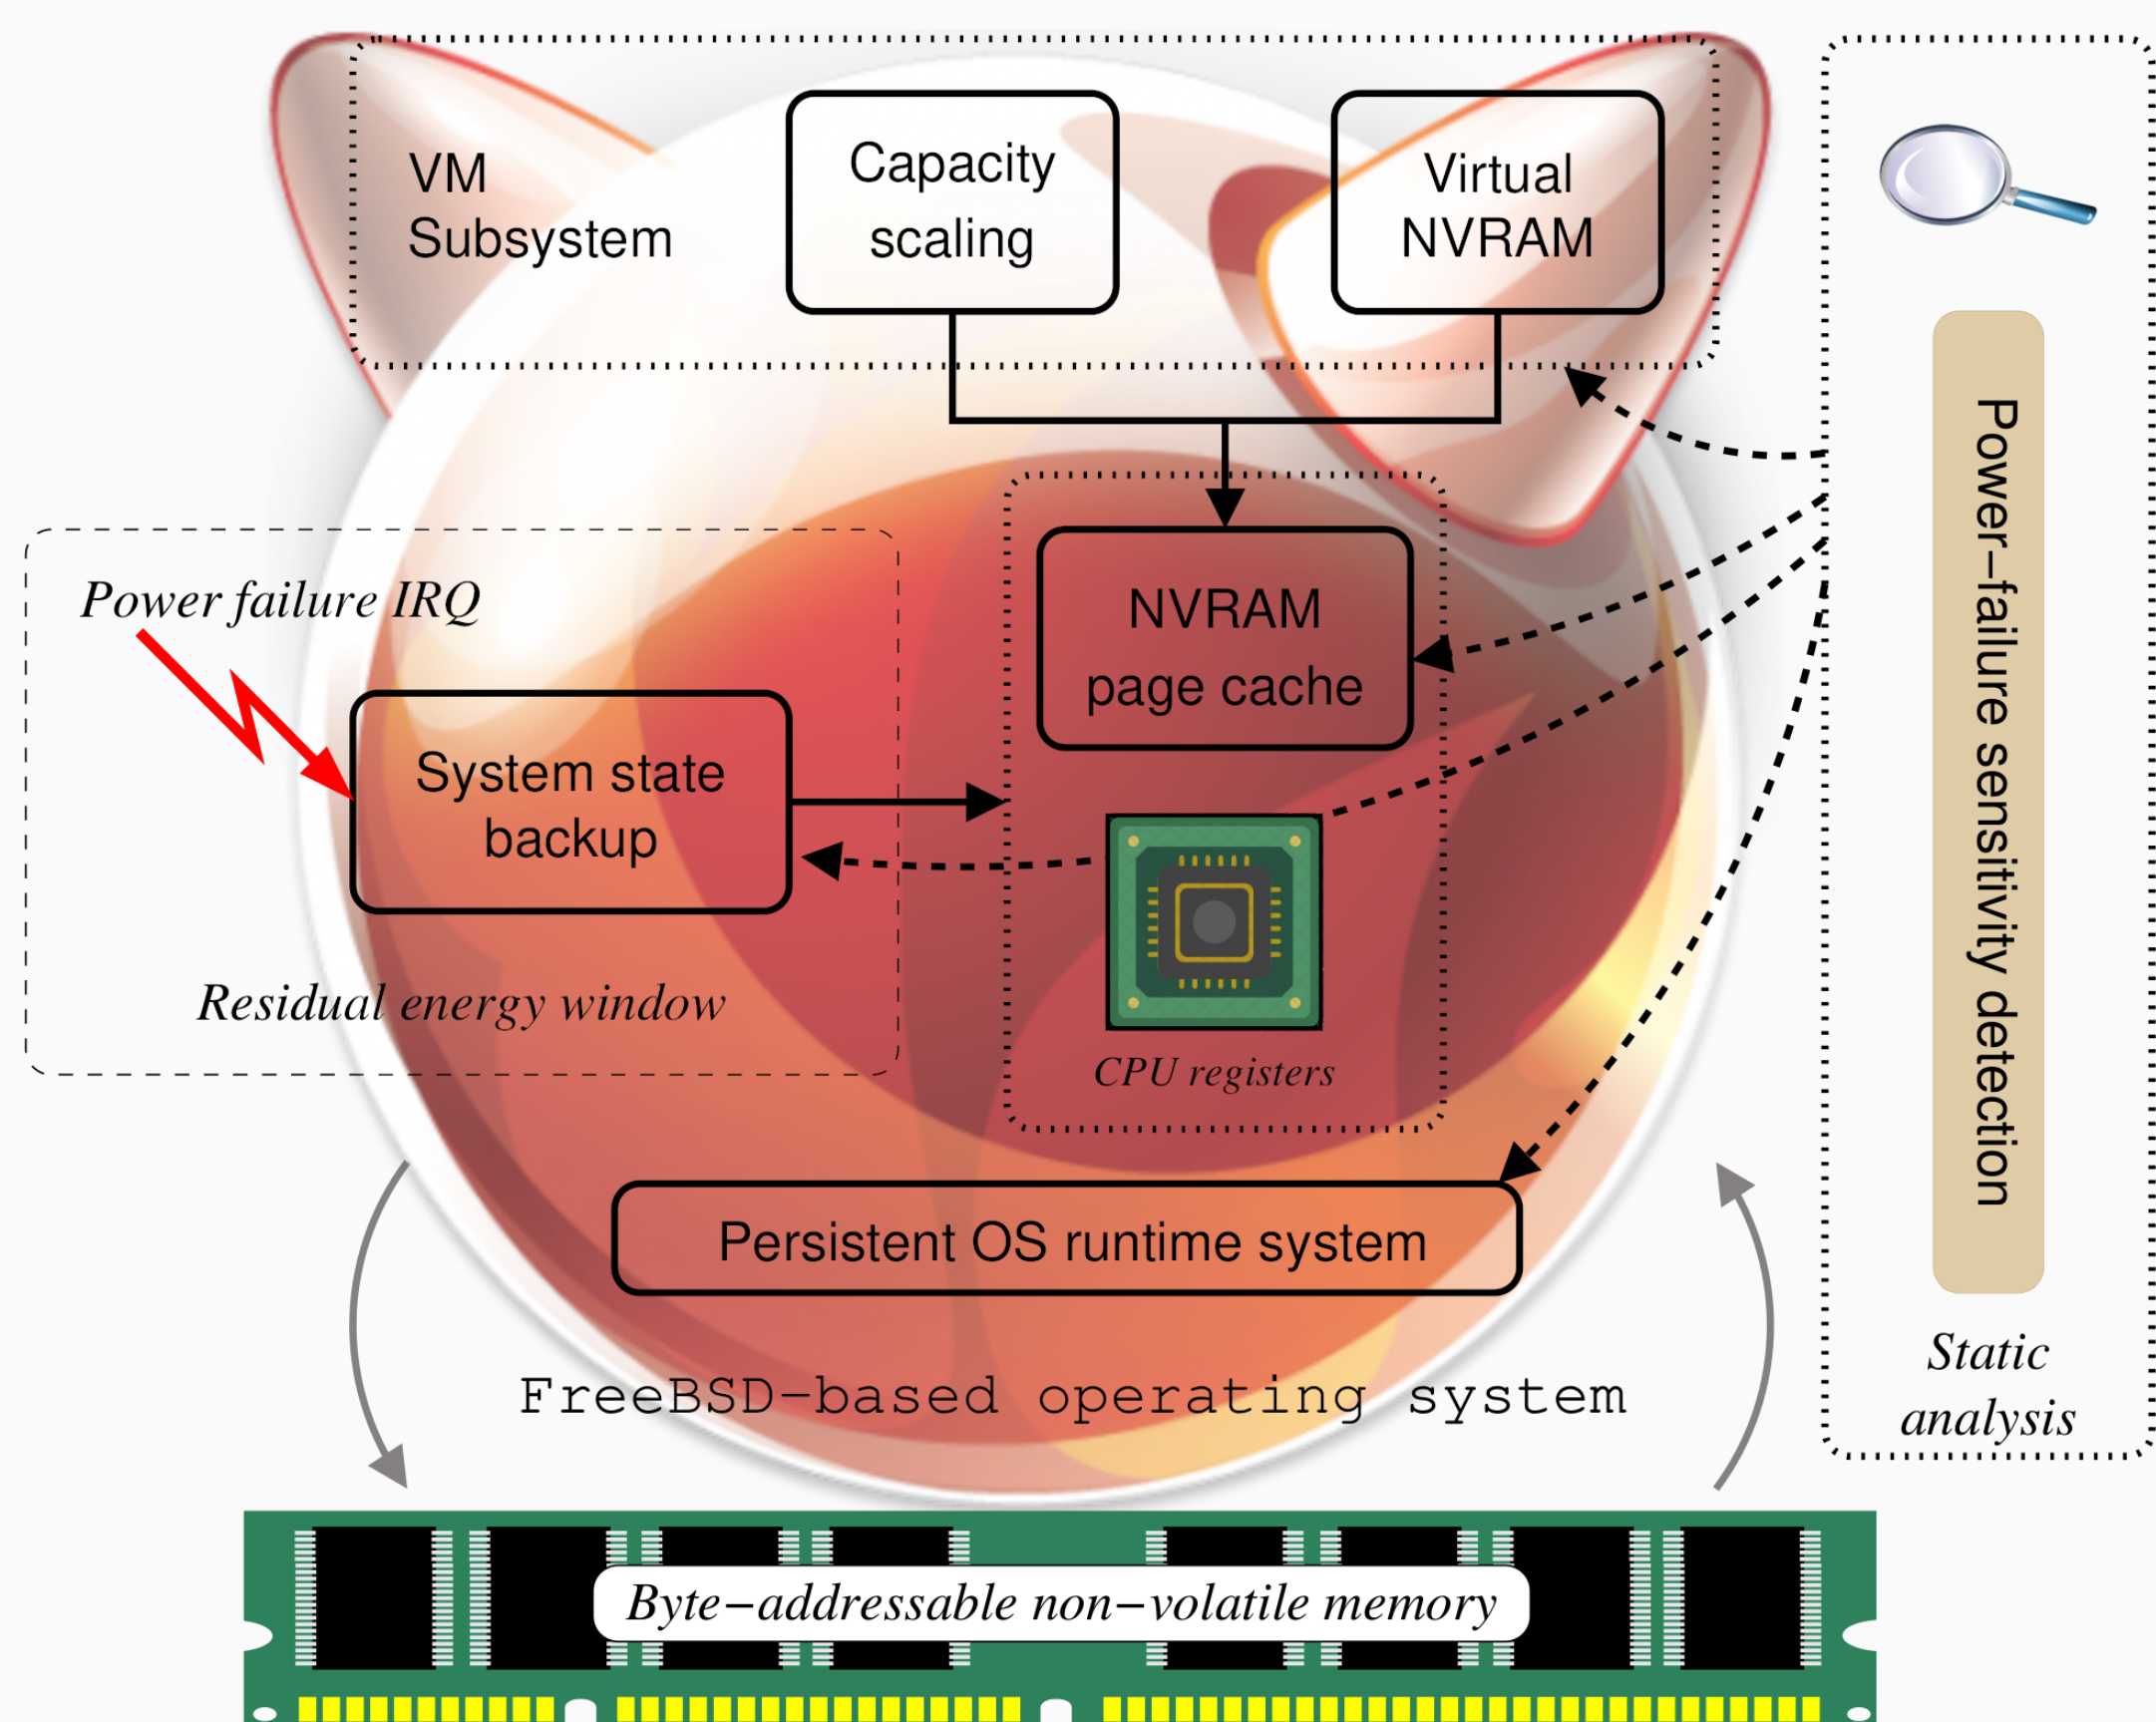
- NVRAM speed is much slower than that of DRAM
- Other storage technologies offer significantly better price/capacity ratios
- Writing NVRAM consumes more power than writing DRAM
- NVRAM requires fail-safe guarantees from the system:
 - power failures on NVRAM-systems may turn a sequential process into a non-sequential one
 - a program has to deal with its own state from previous interrupted runs

Solution Proposal

Integrate these new types of storage-class memory (SCM) appropriately into the memory hierarchy via the virtual-memory subsystem.

- NVRAM-based data cache for conventional storage
 - Size-limited DRAM-based page cache for mutable data in NVRAM
 - Page-cache persistence with strict time guarantees
 - Power failure tolerance: residual-energy-window-compliant page cache
- ⇒ pave the way – for application and system programs – for the use of NVRAM

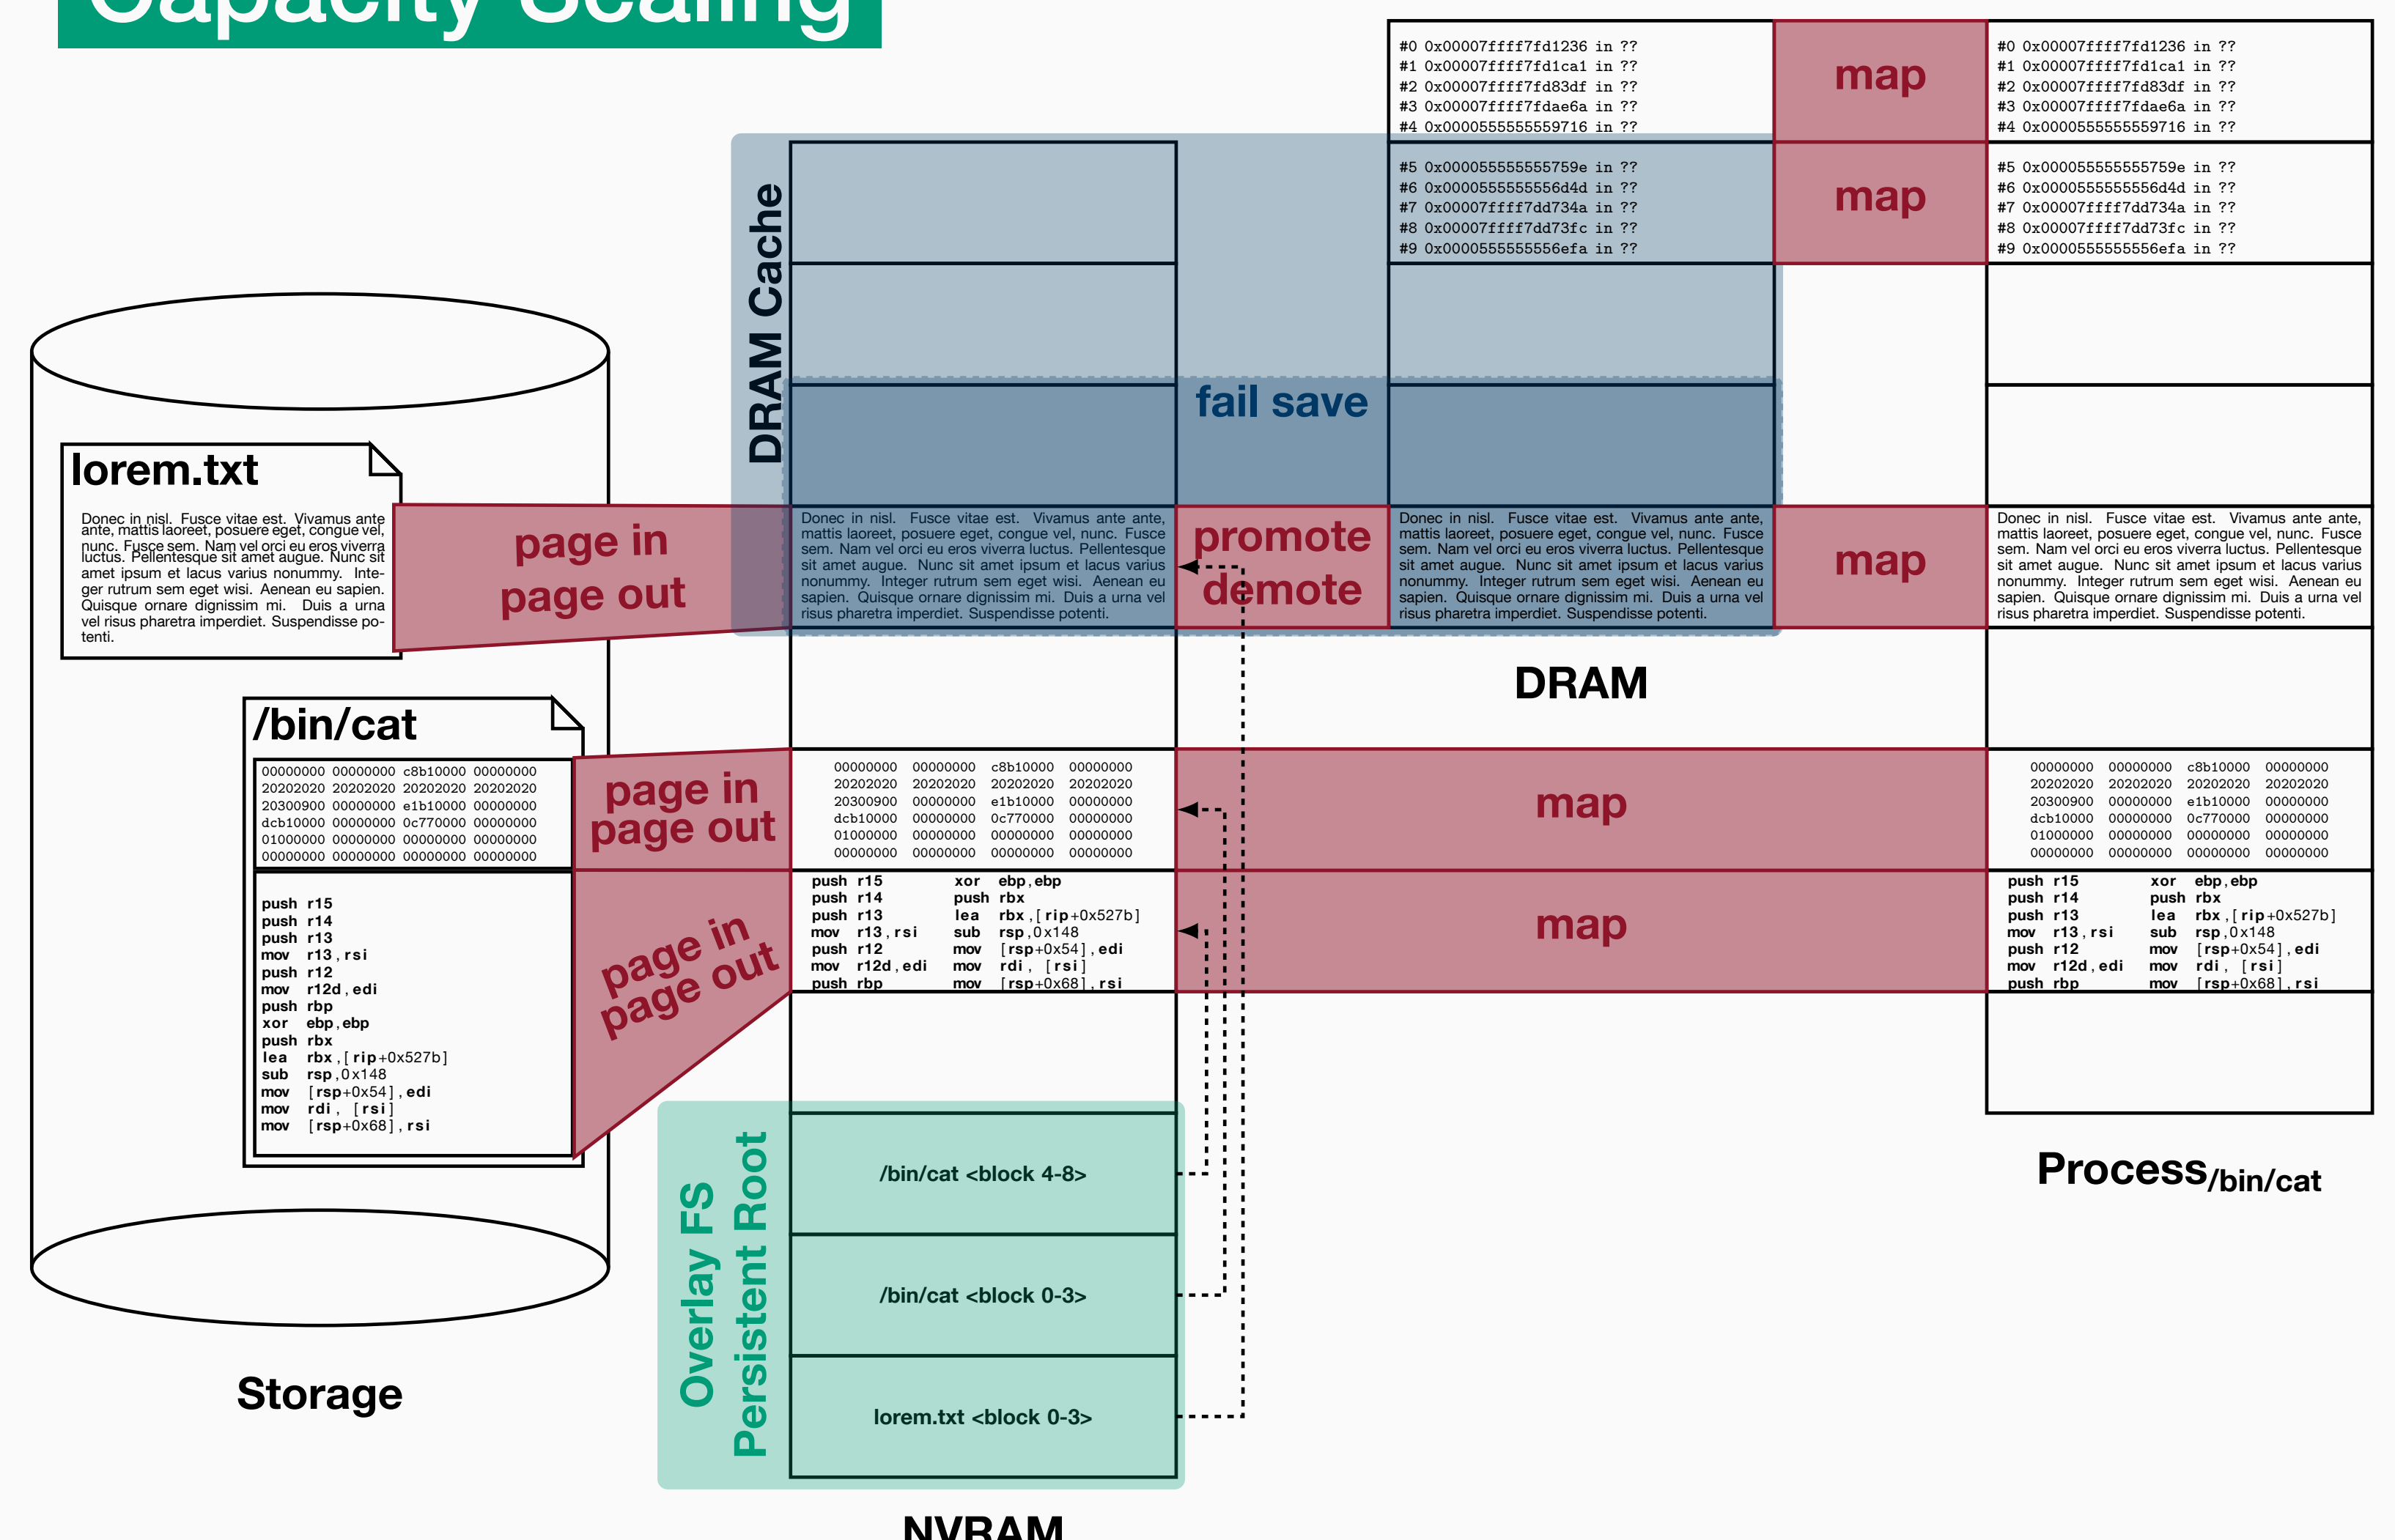
PAVE



Milestones

- NVRAM-based capacity scaling as a kind of OS-managed memory mode
- Extend capacity scaling with persistence guarantees
- Estimate worst-case period for the residual energy from the system power supply
- Limit volatile system state to allow for persistence in case of power failure
- Detection of critical sections with regard to persistence of stale or incomplete state
- Provision of a persistence-aware runtime system for data manipulation

Capacity Scaling



Project Contribution

There are over 35 000 packages available on FreeBSD, which cannot utilize NVRAM:

- Make legacy software benefit from NVRAM
- Consistency without transactional programming
- Capacity scaling