MEXT Performance Benchmarking

How MEXT AI-Powered Predictive Memory™ Works

MEXT continually offloads cold memory pages from DRAM to a 50x lower-cost tier, flash
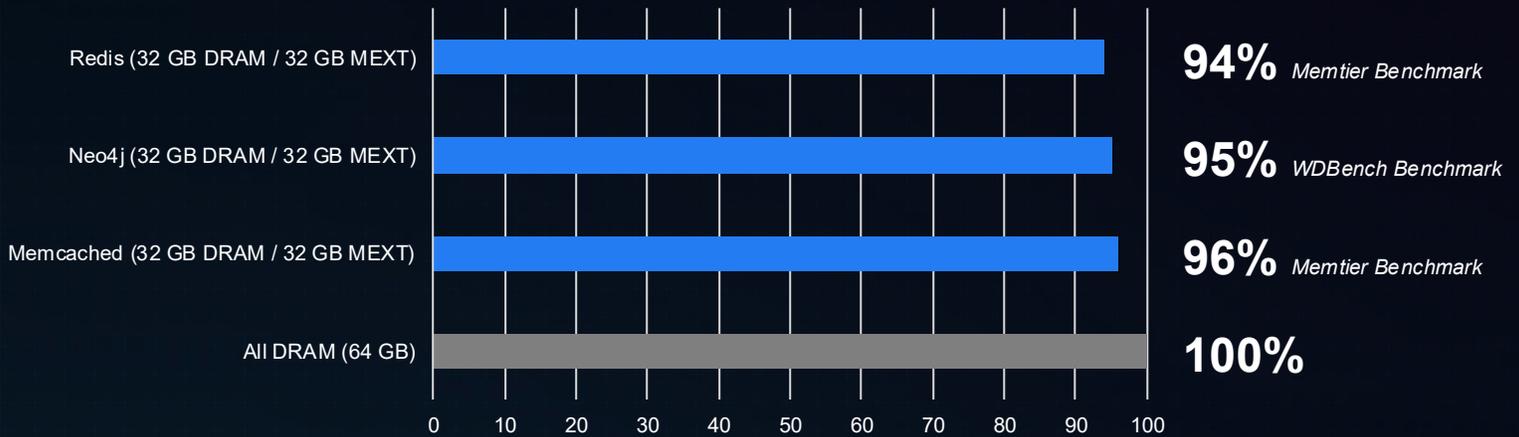
➤

MEXT's AI engine tracks related pages and predicts which pages in flash are likely to get requested by the application

➤

MEXT transparently pushes those pages back into DRAM—even before they are requested by the application

➤

From the application's point of view, all relevant pages can be found in DRAM—maintaining performance within a smaller DRAM footprint

MEXT Performance vs. All DRAM System

With MEXT, customers can achieve comparable application performance using 50% of the DRAM of their original configuration. For instance, when comparing 64 GB DRAM system to a 32 GB DRAM + MEXT system, performance for Redis hits 94% of the all-DRAM configuration, Neo4j hits 95%, and Memcached hits 96%.

	Value	Benchmark
Redis (32 GB DRAM / 32 GB MEXT)	**94%**	*Memtier Benchmark*
Neo4j (32 GB DRAM / 32 GB MEXT)	**95%**	*WDBench Benchmark*
Memcached (32 GB DRAM / 32 GB MEXT)	**96%**	*Memtier Benchmark*
All DRAM (64 GB)	**100%**	

Axis: 0 10 20 30 40 50 60 70 80 90 100

Performance / $ with MEXT

If we look at the amount of performance achieved per dollar spent, we see that leveraging MEXT yields significant efficiencies.

For instance, if we take the Performance / $ of an all-DRAM system to be a reference level of X, MEXT enables a Performance / $ of 1.7X for Redis, 1.7X for Neo4j, and 1.7X for Memcached.

	All DRAM	Redis	Neo4j	Memcached
	1X	**1.7X**	**1.7X**	**1.7X**

Y-axis: 0 25 50 75 100 125 150 175 200