PARTNER WHITE PAPER

When analytics workloads outgrow infrastructure

DCPMM: An alternative to resizing

intel

sas

Rapidly expanding data volumes are testing the limits of even the most advanced analytics infrastructures today. Many clients find out the hard way that when they overcommit system memory (DRAM), performance can take a big hit – fast. In these cases, their data overruns are typically pushed to SWAP space (system and CAS Disk Cache), which is thousands of times slower than DRAM and not well suited to high-performance analytic workloads. Overcommitting system memory can lead to dramatic performance degradation across the board, including analytics and system responsiveness.

Resizing SAS® infrastructure is one proven approach to addressing this issue. But not all clients whose workloads have surpassed their infrastructure are able to resize. Their data may have outgrown their memory capabilities, while their budgets remain fixed – in which case purchasing more servers, cores, or RAM is not an option.

There is no silver bullet – but Intel Optane DC Persistent Memory Modules (DCPMM) can prove to be a valuable alternative to growing the analytics environment in this scenario. Our performance benchmarking of SAS Visual Data Mining and Machine Learning using programmatic workloads (described below) has shown that DCPMM can have a strong positive impact on performance as workload size scales up and reaches capacity, ranging from 4x to 60x faster performance where system Ram is over committed. With DCPMM, it is possible to obtain greater capacity than DRAM for comparable costs, helping future-proof your investments for larger data sets. Alternately, it is possible to achieve reduced costs when compared to purchasing similar capacity in DRAM.

DCPMM requires the latest generation of CPUs (Cascade Lake CPUs). To determine whether your environment specifications meet the minimum requirements (including processor and available DIMM slots) to take advantage of the advances introduced by DCPMM technology, start with a conversation with your Intel representative.

There are also significant cost-benefit reasons to consider DCPMMS, which are significantly less costly than adding DRAM chips, and deliver performance with only marginal degradation, as shown in the results of the procedures we have tested.

To test Optane across a range of SAS procedures, we simulated a financial institution using SAS Visual Data Mining and Machine Learning to predict loan default rates. In order to evaluate creditworthiness in the loan application process, the bank would use analytical models. In this case, an analyst would be using historical data from internal customer transactional systems in tandem with demographic information in order to build predictive models. The internal capacity of the analytics infrastructure would be overloaded by the scale of data required. In this context, we tested six procedures:

- Varreduce
- Treesplit
- Neural Net
- Logselect
- Gradboost
- Forest

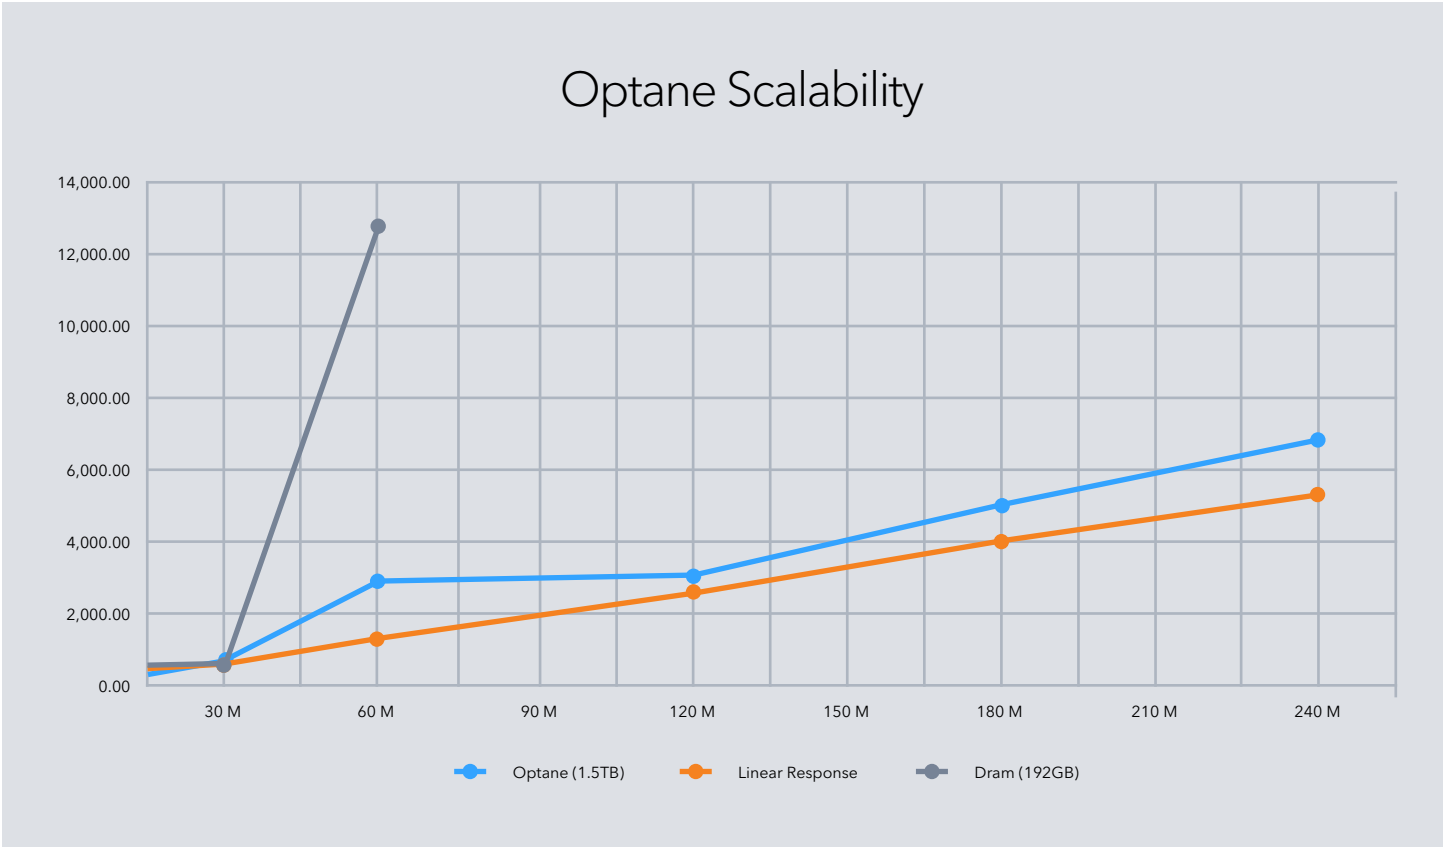

Figure 1. Logistic selection response times comparing DRAM linearity, DCPMM actual and spilling to swap space actuals.

There are also significant cost-benefit reasons to consider DCPMMS, which are significantly less costly than adding DRAM chips, and deliver performance with only marginal degradation, as shown in the results of the procedures we have tested.

Testing Optane: Results

Diagram 1 shows a typical graph of the difference in performance between Intel Optane DCPMM and system DRAM – in this case, the SAS Logistic procedure is represented. Degradation occurs quickly and dramatically when using DRAM – the gray vertical line. But Optane supports near-DRAM response times (blue line) that are comparable to linear projections (orange line) resizing.

Here's a more detailed look at the difference between processing times in DRAM compared to those using Intel Optane. In every case, DCPMM modules deliver considerably faster processing times – in the case of Gradboost, it is more than 60 times faster than DRAM.

A Practical Alternative

Is Optane the right option for clients who have reached the limits of their current SAS infrastructure with growing data management and analytics demands? Maybe. There are inherent tradeoffs in pursuing a DCPMM-based approach to expanding processing capacity. Regardless, Optane warrants careful consideration along with resizing infrastructure as a practical, budget-conscious approach.

Benchmark Environment System Details

Lenovo: Think System SR650

- CPU: Intel Xeon Platinum 8268 CPU @ 2.90GHz (2 sockets)
- Cores: 48 total (96 threads)
- RAM: DDR4/184GB (12 x 16GB), 2666Mhz
- Intel Optane: 1.5TB (12 x 128GB)

CAS Disk Cache:

- 3x -1.6TB NVMe SSD

Software:

- Red Hat Enterprises Linux Server release 7.6
- SAS Visual Data Mining and Machine Learning 8.4 (19w21 + updates, programmatic only)

SAS Procedure	DRAM Response (hh:mm:ss)	DCPM Response (hh:mm:ss)	Ratio (times faster)
Variable Reduction	00:40:15	00:03:19	12x
Decision Tree	04:53:43	00:09:49	30x
Neural Network	05:52:24	01:15:22	4.6x
Logistic	03:35:25	00:46:57	4.5x
Gradient Boosting	59:41:33	00:55:38	63x
Forest	01:46:46	00:21:48	5x

Table 1. Represents response time comparisons between DRAM and DCPMM in an overcomitted scenario (237GB input table loaded into 192GB DRAM) resulting in active paging without DCPMM. Adding DCPMM eliminates paging, enhancing performance by factors between 4x and 60x.