PrimoCache

The Detailed Performance Evaluation of Intel Optane SSD 900P and PrimoCache and the Application of Both in Server/Workstation

Doc. No.	:	RS-SW-PCC-50-18-01			Ver. 1.2
Creation Date	:	2018-03-09	Last Mod Date	:	2019-04-20
Status	:	Final	Classification	:	Public

OVERVIEW

Intel's Optane series products, based on 3D XPoint technology, has been attracting attention since its launch with its ultra-high performance, ultra-low latency and long life. So far the published products include High-End P4800X series for data centers and the Optane Memory series for Desktop Cache Acceleration, the SSD 900P Series and 800P Series for High-End Desktop, Workstation and Storage Acceleration.

According to Intel's official specifications, the SSD 900P series comes in two form factors: a HHHL PCIe Add-In Card and a 2.5" U.2 drive, both with a PCIe 3.0 x4 interface. Currently, it supports capacity of 280GB and 480GB, and will be followed later by higher-capacity models. It claims impressive performance, up to 2,500MB/s and 2,000MB/s in sequential reads and writes, respectively, and up to 550,000 IOPS and 500,000 IOPS in 4KB random reads and writes, as well as low to 10µs in access latency. The endurance-per-GB is 18.69TBW, which comes out to be 10 complete drive writes per day over the course of the 5-year warranty.

Although Intel launched the SSD 900P series is officially targeting the consumer market, the SSD 900P performs exceptionally well in terms of read and write performance, latency, and IO throughput, making it suitable for a variety of workstations and small and medium-sized servers. Especially when paired with PrimoCache software, the SSD 900P can serve as a fast cache for other slow hard drives, greatly improving the read/write speed and IO throughput of these hard drives, even closing to or reaching the performance of the SSD 900P. PrimoCache + SSD 900P cache + large-capacity slow HDD/SSD is a cost-effective solution for users who do not have enough budget to purchase expensive large-capacity and high-performance SSD. On the other hand, this caching solution does not require any changes to the existing hardware configuration and software environment of workstations or servers, and does not require any data migration. It only requires the installation of SSD 900P and PrimoCache software to bring about high performance improvement. It's quite simple and convenient, and it is ideal for users who need to improve performance but not modify existing hardware or software.

This article will evaluate the original read and write performance of the SSD 900P and its performance as the PrimoCache cache in detail. This article does not show you how to use the PrimoCache software or how to configure the SSD 900P as a cache. For instructions on these, please read the PrimoCache help documentation or Quick Start Guide. The PrimoCache software can be downloaded from the official website.

TEST SCHEME

The Optane product tested in this article is an Intel Optane SSD 900P add-in card. The test platform and system software environment are shown in the following table.

Baseboard	Intel S2600CW
CPU	Intel Xeon CPU E5-2698 v3 @ 2.30GHz x1
RAM	Samsung 16GB DDR4-2133 RDIMM x4, total 64GB
Drives	Seagate Constellation ES ST1000NM0011 1TB 3.5" Drive (7200RPM, SATA3, 64MB)
	Intel Optane SSD 900P 280GB PCI-E HHHL
OS	Windows Server 2012 Standard v6.2.9200
NVMe Driver	Intel NVMe Driver 3.2.0.1002
Test Tool	Microsoft Diskspd v2.0.17
PrimoCache	Server Edition v3.0.2 Beta

The PrimoCache cache configuration used during the test is shown in the figure below. Only level-2 (SSD) caching is deployed, no level-1 (RAM) caching.

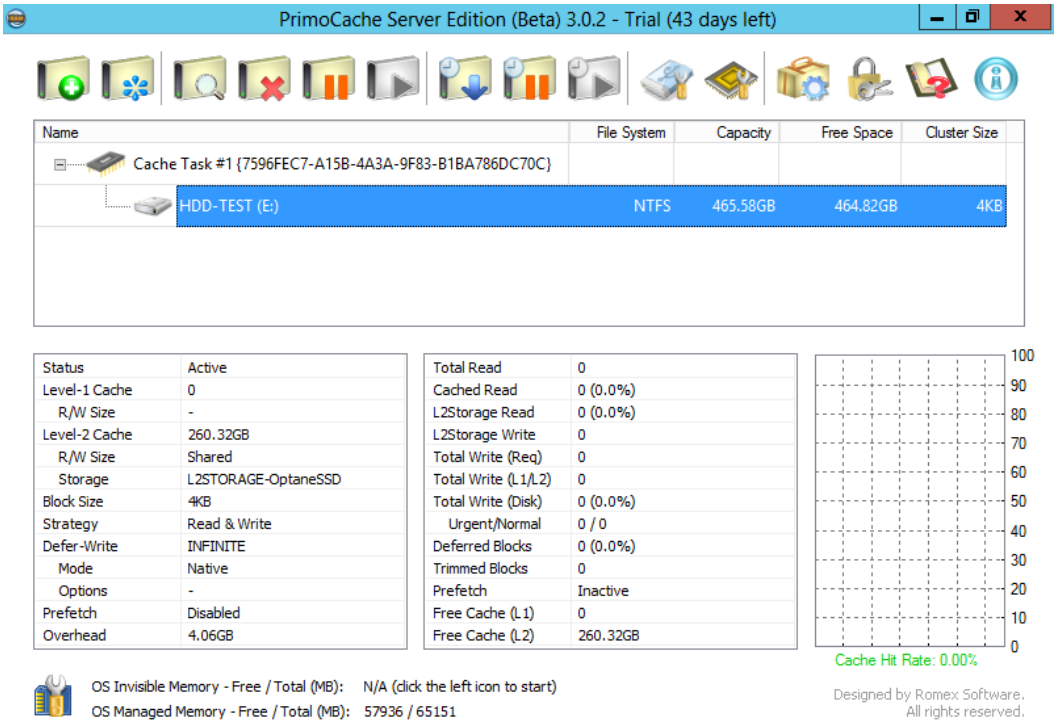

The test items include sequential standard read/write, random standard read/write, sequential mixed read-write, and random mixed read-write. The test data block size covers a full range of 4KB to 1MB. The test load takes into account the low/middle/high demand situations, including single-threaded and multi-threaded tests at different queue depth (QD), up to 32-thread and 64-QD. Test scenarios cover full and empty drives, whole and partial storage space, as well as a write-life impact test. The overall test scheme reflects the read/write speed and IO performance of the test object in various possible scenarios, as expected to be found in the practical application environment.

TEST REPORT SPECIFICATION

All test results reported in this article show the measured results of each test in a chart. Each row in each report contains four charts, showing the four measured values of a set of tests in turn: data transfer rate (MB/s), IOPS, average latency (milliseconds) and average CPU utilization (%).

In the test reports of standard read/write tests and mixed read/write tests, each chart typically displays three broken lines, labeled OptaneSSD 900P, ST1000NM0011, and PrimoCache L2. OptaneSSD 900P and ST1000NM0011 refer to the results of the Intel Optane SSD 900P 280GB and the Seagate ST1000NM0011 drive test without the installation of PrimoCache software. PrimoCache L2 is the result of testing on the ST1000NM0011 drive with the SSD 900P served as its cache.

The QxTy in test reports indicates that the test is operated at the same time by y threads while each thread has a queue depth of x, which comes out to be x*y IO operations executed simultaneously. For example, Q1T1 indicates single thread single queue depth, i.e. single IO operation, while Q64T32 indicates 32 threads 64 queue depth, i.e. 2048 IO operations at the same time.

Due to the large number of test results, only a small number of typical charts are cited in this article for analysis. To see all the detailed data and charts, please download the attachment at the end of the article.

STANDARD READ/WRITE TEST

The standard read/write test mainly tests the performance of sequential standard read, sequential standard write, random standard read, and random standard write. The following charts show the results of the complete-drive test. Complete- drive test refers to read/write tests on the whole space of the entire storage device. Before the test starts, the whole drive has been filled by test data. Therefore, the complete-drive test here is also a full-drive test.

SEQUENTIAL STANDARD READ

Sequential standard read usually focuses on the processing capability of 64KB to 1MB data blocks. As can be seen from the chart below, the sequential read speed of SSD 900P in Q1T1 which is under extremely low workload is about 1,700MB/s (64KB) to 2,300MB/s (1MB), while in multi-thread or multi-QD the sequential read speed of 64KB to 1MB is basically stable at the Intel claimed rate of 2,500MB/s, **even close to 2,600MB/s**! In the practical application of server/workstation, most of the storage devices are working in multi-threaded and multi-QD states, so the SSD 900P can take full advantage of its maximum performance.

At the same time, it can be seen that after the SSD 900P is used as the cache for the ST1000NM0011 drive, the sequential read speed of the ST1000NM0011 approaches or reaches the SSD 900P's performance. Under certain circumstances, PrimoCache can fully exert the SSD 900P performance, making the performance of cached drive even beyond the SSD 900P itself.

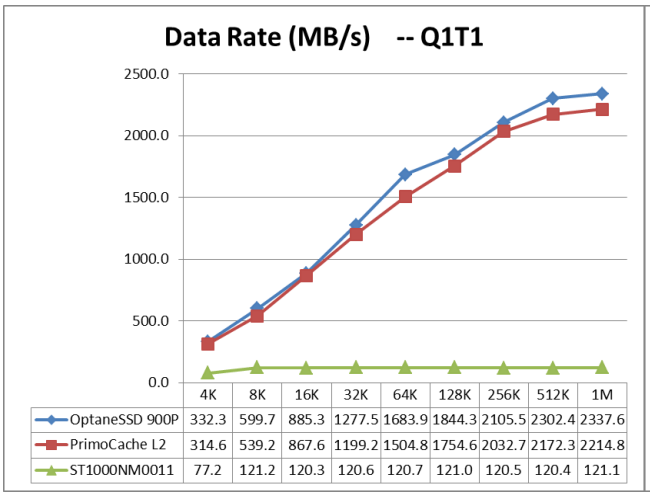

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	332.3	599.7	885.3	1277.5	1683.9	1844.3	2105.5	2302.4	2337.6
PrimoCache L2	314.6	539.2	867.6	1199.2	1504.8	1754.6	2032.7	2172.3	2214.8
ST1000NM0011	77.2	121.2	120.3	120.6	120.7	121.0	120.5	120.4	121.1

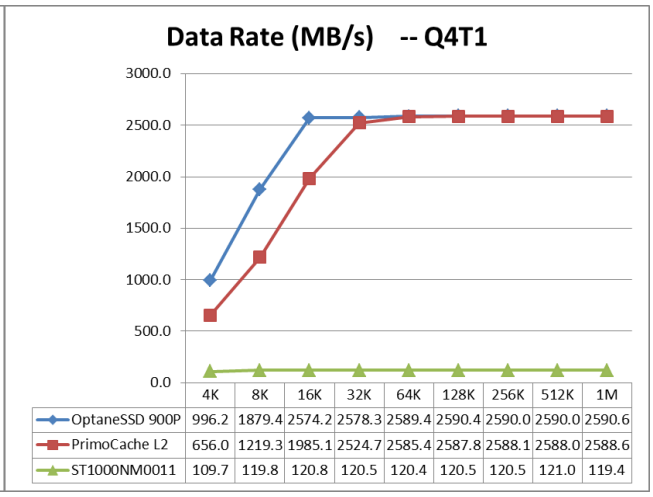

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	996.2	1879.4	2574.2	2578.3	2589.4	2590.4	2590.0	2590.0	2590.6
PrimoCache L2	656.0	1219.3	1985.1	2524.7	2585.4	2587.8	2588.1	2588.0	2588.6
ST1000NM0011	109.7	119.8	120.8	120.5	120.4	120.5	120.5	121.0	119.4

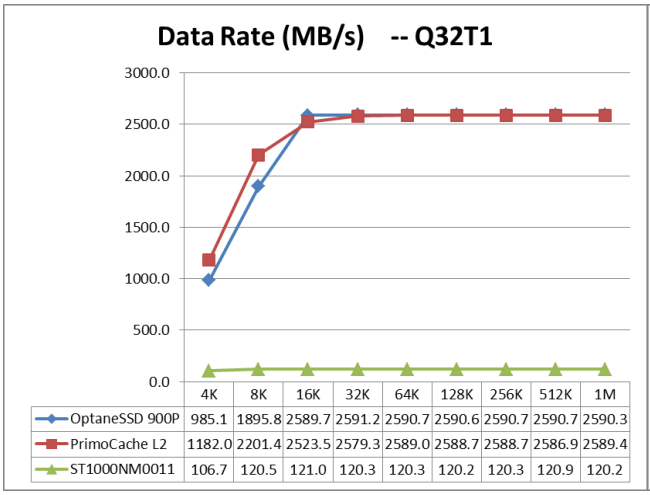

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	985.1	1895.8	2589.7	2591.2	2590.7	2590.6	2590.7	2590.7	2590.3
PrimoCache L2	1182.0	2201.4	2523.5	2579.3	2589.0	2588.7	2588.7	2586.9	2589.4
ST1000NM0011	106.7	120.5	121.0	120.3	120.3	120.2	120.3	120.9	120.2

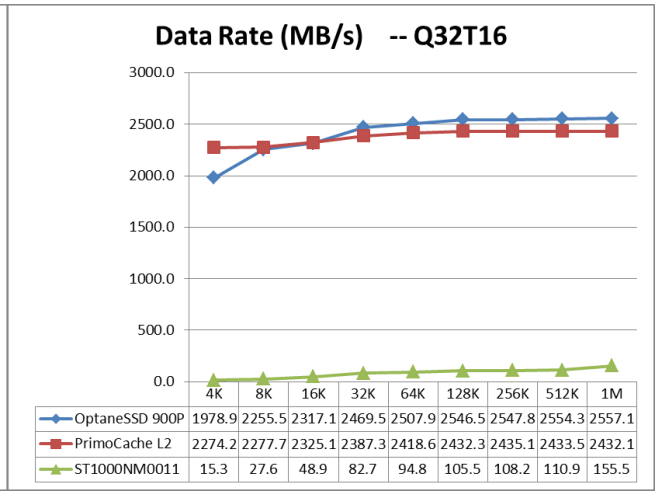

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	1978.9	2255.5	2317.1	2469.5	2507.9	2546.5	2547.8	2554.3	2557.1
PrimoCache L2	2274.2	2277.7	2325.1	2387.3	2418.6	2432.3	2435.1	2433.5	2432.1
ST1000NM0011	15.3	27.6	48.9	82.7	94.8	105.5	108.2	110.9	155.5

SEQUENTIAL STANDARD WRITE

Sequential standard write, similar to sequential standard read, also focuses on the processing capability of 64KB to 1MB data blocks. From the chart, we can see that the sequential write speed of the SSD 900P is about 1,400MB/s (64KB) to 1,900MB/s (1MB) in Q1T1, while in multi-QD or multi-thread (less than 32 threads) the sequential write speed of 64KB to 1MB is basically stable at the claimed rate of about 2,000MB/s, **the highest even close to 2,200MB/s**.

Similarly, the sequential write performance of the PrimoCache cache is also close to or reaches the performance of the SSD 900P. In addition, the test results also show an interesting phenomenon. With the increase of the number of simultaneous threads, the maximum sequential write speed of the SSD 900P is reduced from approximately 2,200MB/s (single thread) to approximately 1,800MB/s (32 threads), however, the PrimoCache cache can still be maintained at 2,000MB/s even under 32 threads, with the maximum utilization of SSD 900P performance, making its performance surpass the SSD 900P. One possible reason for this phenomenon is that the SSD 900P controller chip is limited by its own hardware resources and cannot fully exploit the maximum performance of the storage medium, whereas the PrimoCache cache can allocate the hardware resources of the entire computer system to achieve maximum performance.

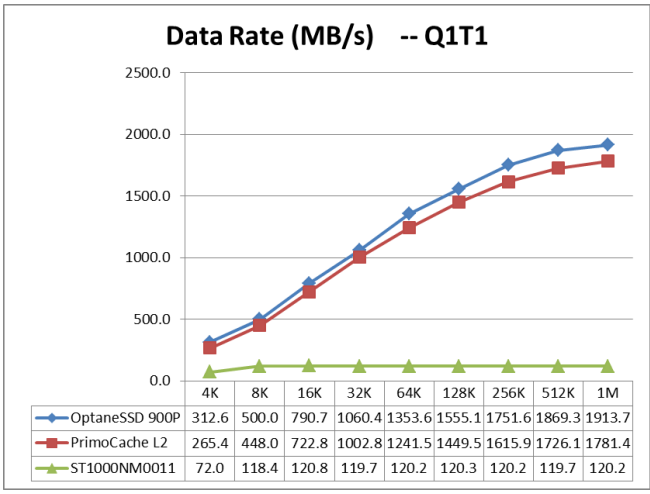

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	312.6	500.0	790.7	1060.4	1353.6	1555.1	1751.6	1869.3	1913.7
PrimoCache L2	265.4	448.0	722.8	1002.8	1241.5	1449.5	1615.9	1726.1	1781.4
ST1000NM0011	72.0	118.4	120.8	119.7	120.2	120.3	120.2	119.7	120.2

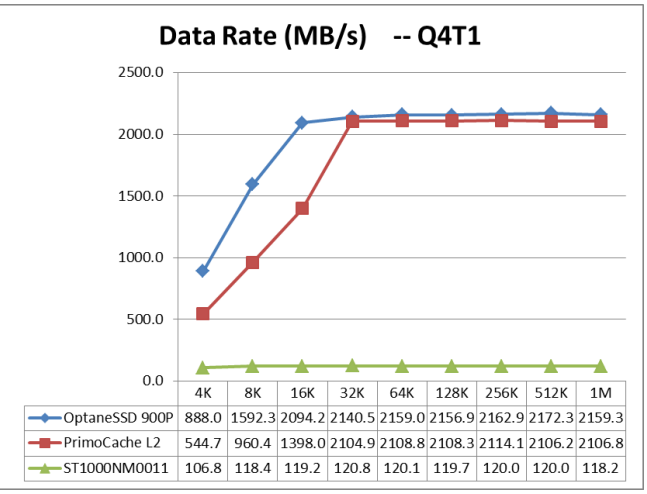

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	888.0	1592.3	2094.2	2140.5	2159.0	2156.9	2162.9	2172.3	2159.3
PrimoCache L2	544.7	960.4	1398.0	2104.9	2108.8	2108.3	2114.1	2106.2	2106.8
ST1000NM0011	106.8	118.4	119.2	120.8	120.1	119.7	120.0	120.0	118.2

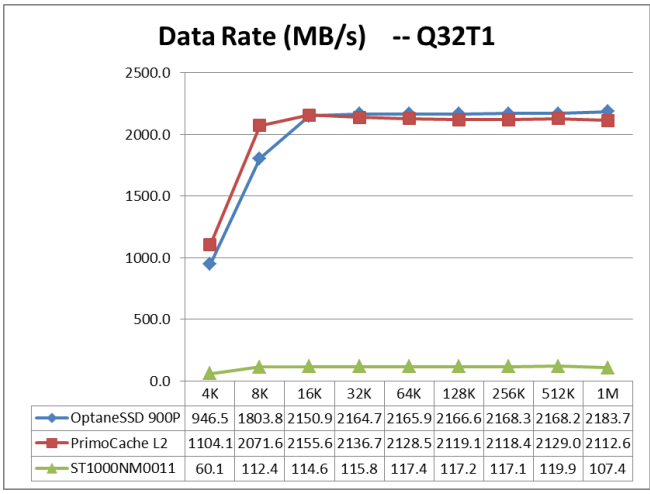

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	946.5	1803.8	2150.9	2164.7	2165.9	2166.6	2168.3	2168.2	2183.7
PrimoCache L2	1104.1	2071.6	2155.6	2136.7	2128.5	2119.1	2118.4	2129.0	2112.6
ST1000NM0011	60.1	112.4	114.6	115.8	117.4	117.2	117.1	119.9	107.4

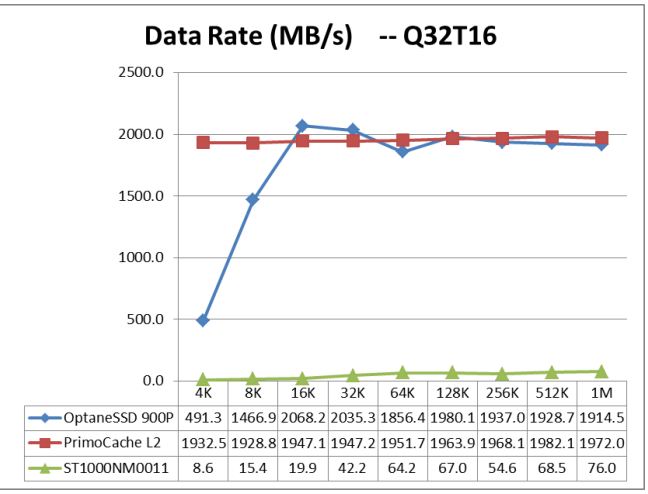

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	491.3	1466.9	2068.2	2035.3	1856.4	1980.1	1937.0	1928.7	1914.5
PrimoCache L2	1932.5	1928.8	1947.1	1947.2	1951.7	1963.9	1968.1	1982.1	1972.0
ST1000NM0011	8.6	15.4	19.9	42.2	64.2	67.0	54.6	68.5	76.0

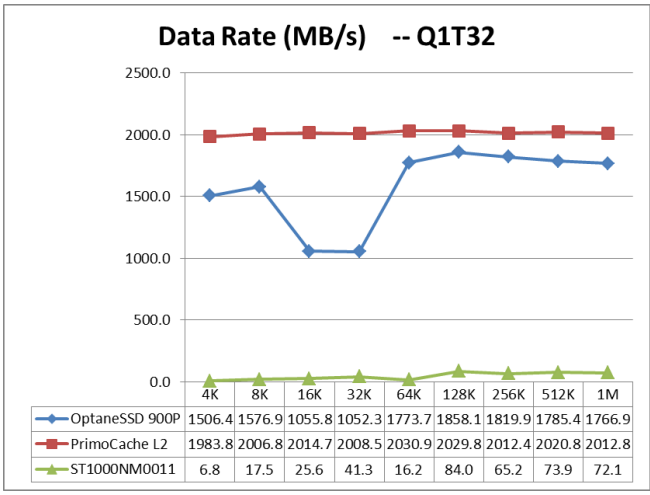

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	1506.4	1576.9	1055.8	1052.3	1773.7	1858.1	1819.9	1785.4	1766.9
PrimoCache L2	1983.8	2006.8	2014.7	2008.5	2030.9	2029.8	2012.4	2020.8	2012.8
ST1000NM0011	6.8	17.5	25.6	41.3	16.2	84.0	65.2	73.9	72.1

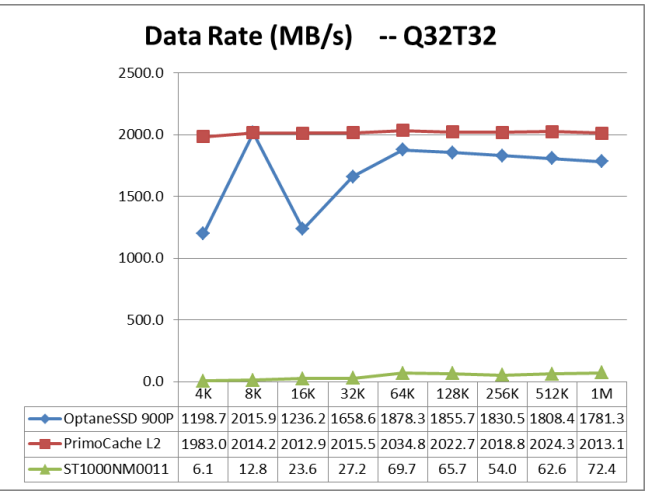

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	1198.7	2015.9	1236.2	1658.6	1878.3	1855.7	1830.5	1808.4	1781.3
PrimoCache L2	1983.0	2014.2	2012.9	2015.5	2034.8	2022.7	2018.8	2024.3	2013.1
ST1000NM0011	6.1	12.8	23.6	27.2	69.7	65.7	54.0	62.6	72.4

RANDOM STANDARD READ

For the random standard read, we generally focus on the performance of 4KB to 64KB data blocks. Here we only focus on the IOPS of the 4KB block. From the chart, we can see that the IOPS value of 4KB random read of the SSD 900P in Q1T1 is close to 90,000. With the increase of the number of threads or queue depth, the IOPS value is also significantly increased. For e.g. in Q4T8, IOPS is already nearing

580,000. The test results show that **the highest IOPS value of 4KB random read is stable at nearly 590,000, which is even higher than the official rate of 550,000**.

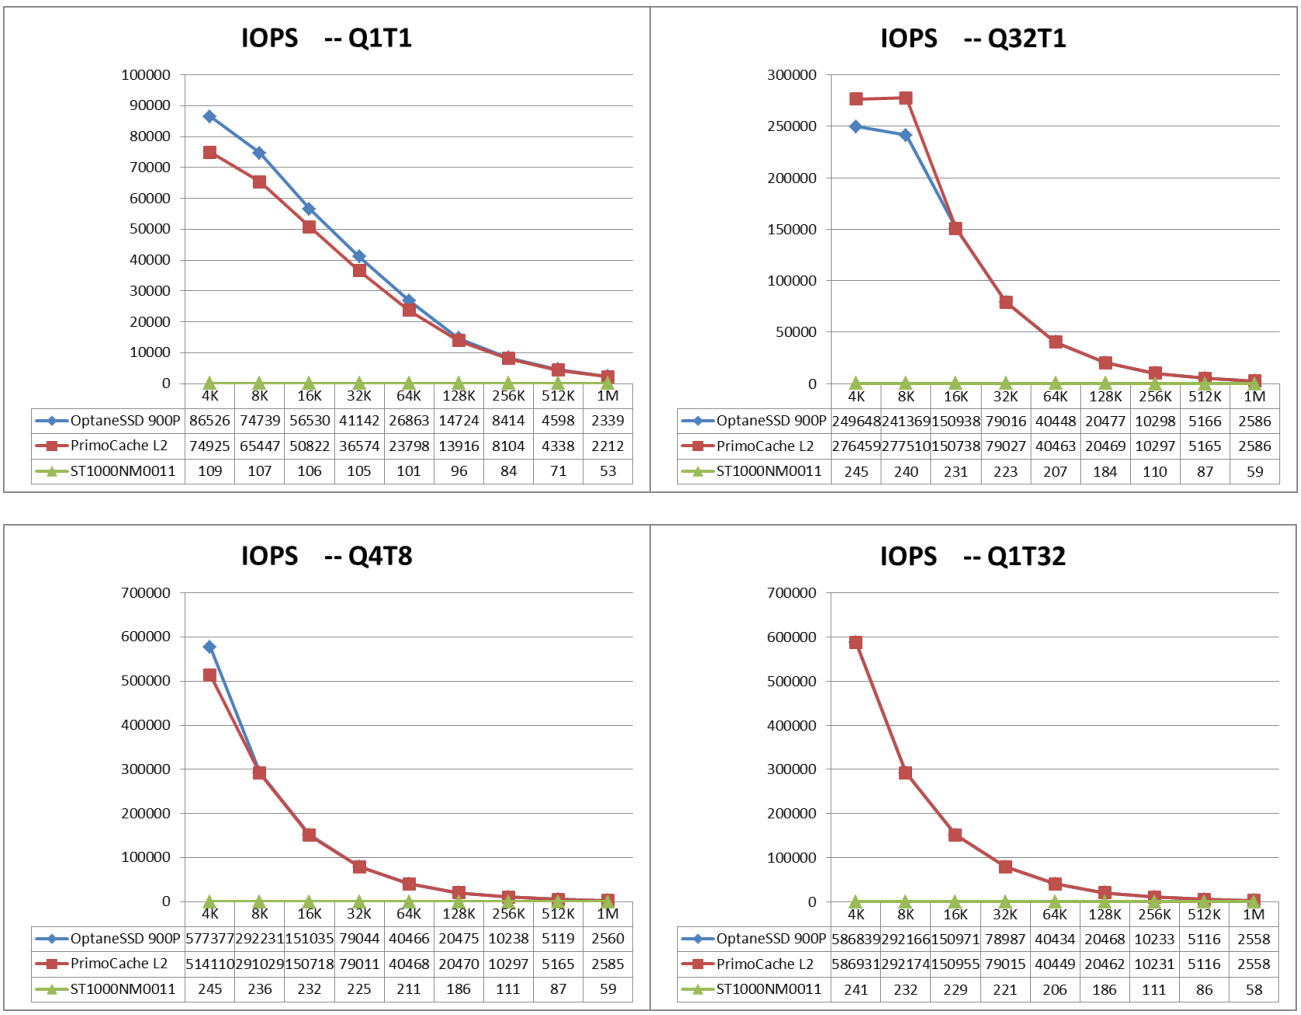

The PrimoCache cache for the performance of the 4KB block will be weaker than the SSD 900P in most cases. This is because when the cache handles I/O requests, it will take additional process time for caching works besides delivering requests to the SSD 900P. Though this extra process time is very little, it still forms a noticeable proportion of the total process time, as the SSD 900P also takes quite little time to complete an I/O request. The difference in processing time reflected in the IOPS is more obvious, especially in the case of low thread count and low queue depth. However, in the case of low thread count and high queue depth (such as Q32T1), due to the PrimoCache cache can fully utilize the performance of the SSD 900P, the IOPS value of 4KB is higher than that of the SSD 900P.

From the latency measurement results in Q1T1, we can see that the minimum average latency of the SSD 900P 4KB random read is 11µs, which is basically consistent with the Intel claimed 10µs. The corresponding latency of the PrimoCache cache is 13µs, which is 2µs more than the SSD 900P.

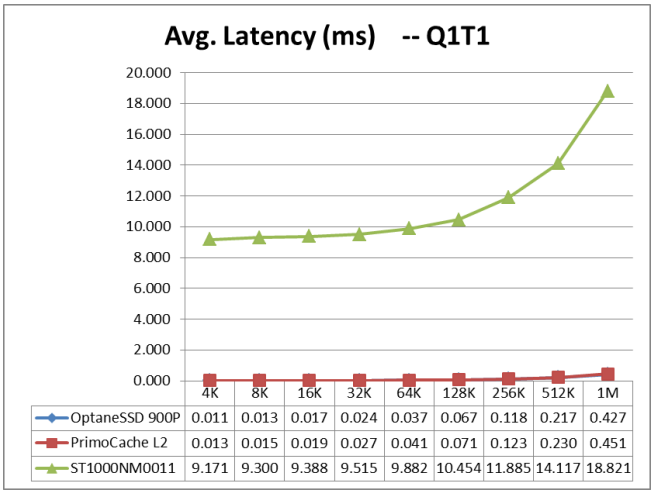

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	0.011	0.013	0.017	0.024	0.037	0.067	0.118	0.217	0.427
PrimoCache L2	0.013	0.015	0.019	0.027	0.041	0.071	0.123	0.230	0.451
ST1000NM0011	9.171	9.300	9.388	9.515	9.882	10.454	11.885	14.117	18.821

RANDOM STANDARD WRITE

The test results of the random standard write are similar to those of random read, the IOPS value of 4KB random write of SSD 900P in Q1T1 is close to 80,000. With the increase of the number of threads or queue depth, the IOPS value increases significantly. **The highest can reach more than 550,000, higher than the official 500,000.**

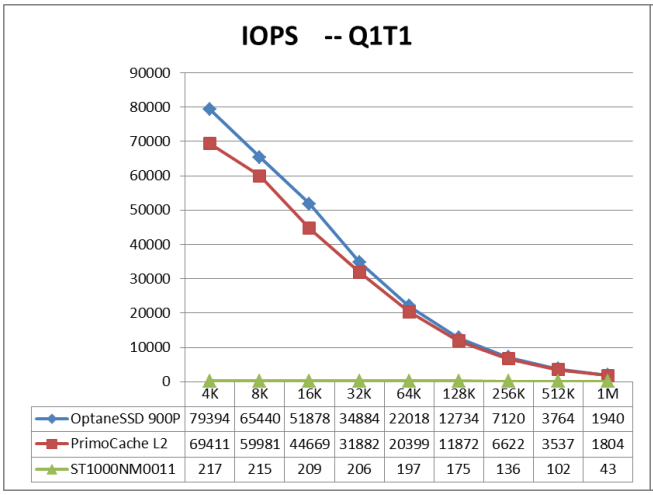

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	79394	65440	51878	34884	22018	12734	7120	3764	1940
PrimoCache L2	69411	59981	44669	31882	20399	11872	6622	3537	1804
ST1000NM0011	217	215	209	206	197	175	136	102	43

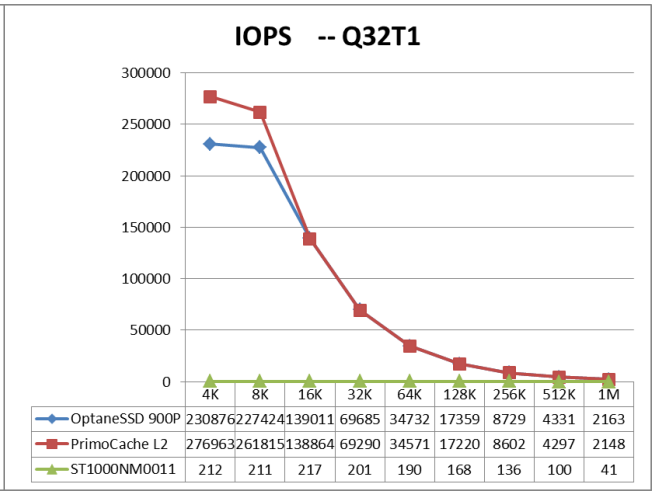

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	230876	227424	139011	69685	34732	17359	8729	4331	2163
PrimoCache L2	276963	261815	138864	69290	34571	17220	8602	4297	2148
ST1000NM0011	212	211	217	201	190	168	136	100	41

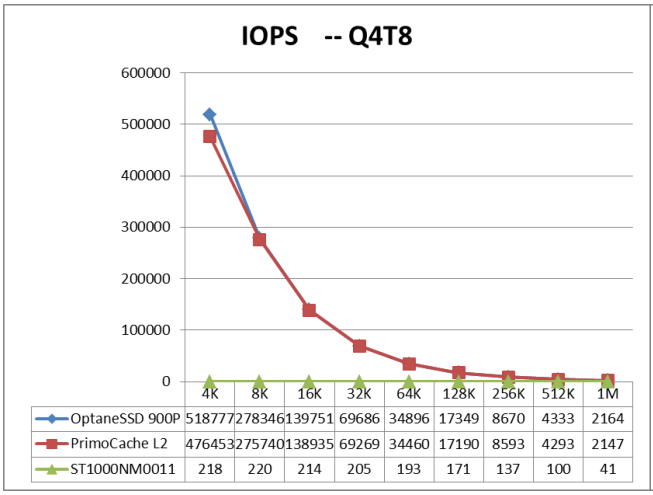

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	518777	278346	139751	69686	34896	17349	8670	4333	2164
PrimoCache L2	476453	275740	138935	69269	34460	17190	8593	4293	2147
ST1000NM0011	218	220	214	205	193	171	137	100	41

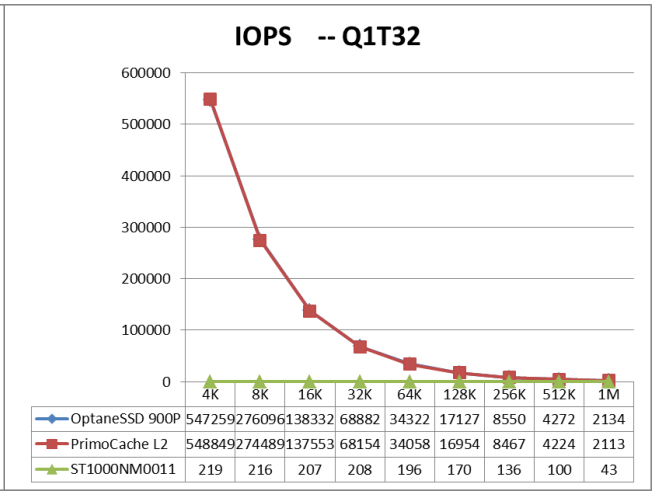

	4K	8K	16K	32K	64K	128K	256K	512K	1M
OptaneSSD 900P	547259	276096	138332	68882	34322	17127	8550	4272	2134
PrimoCache L2	548849	274489	137553	68154	34058	16954	8467	4224	2113
ST1000NM0011	219	216	207	208	196	170	136	100	43

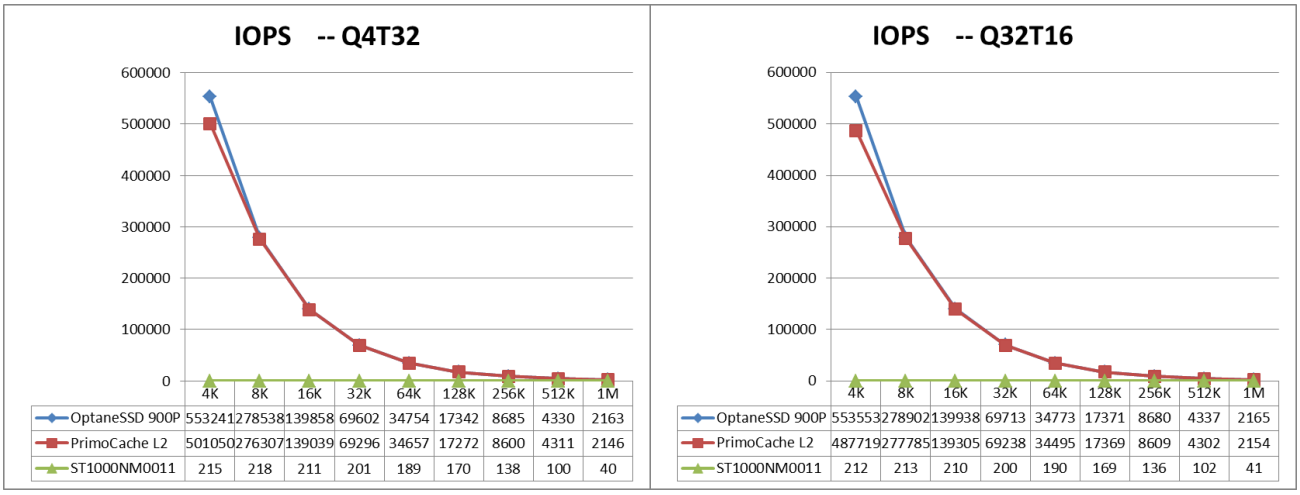

From the latency measurement results in Q1T1, we can see that the minimum average latency of the SSD 900P 4KB random write is 12µs, which is slightly higher than the official 10µs. The corresponding latency of the PrimoCache cache is 14µs, which is also 2µs more processing overhead than the SSD 900P, same as the result of 4KB random read.

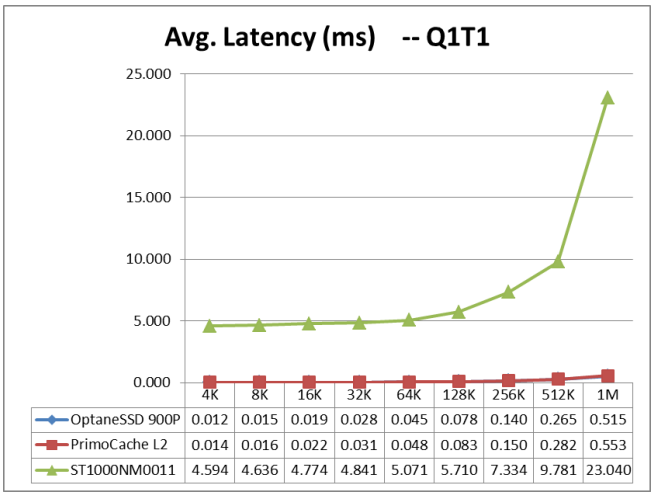

MIXED READ-WRITE TEST

The mixed read-write test is mainly to test the performance when the storage device reads and writes at the same time. In this article, we tested the performance of sequential and random mixed read-write with different data block sizes under different thread counts and different queue depth. Typical test result charts are selected here for illustration. The x-coordinate in the chart represents "% write", that is, the percentage of write IOs that accounts for the total number of IOs in the test. For example, "30% write" means that test IOs are composed of 30% write IOs and 70% read IOs. "0% write" is complete read (i.e., standard read) and "100% write" is complete write (i.e., standard write).

SEQUENTIAL MIXED READ-WRITE

The following chart shows the test results of sequential mixed read-write of the 128KB block in Q1T1, Q32T1, Q1T16 and Q8T8. It can be seen that with the change of read/write ratio, SSD 900P presents a corresponding nearly linear change of 128KB sequential mixed read-write performance in proportion

to the full read/write performance. This is an amazing result, as many other SSDs or mechanical hard drives have the phenomenon of high full read/write performance but a steep drop in mixed read-write performance.

Similar to the standard read/write test results, the performance of the PrimoCache cache in sequential mixed read-write is basically close to or up to the performance of the SSD 900P.

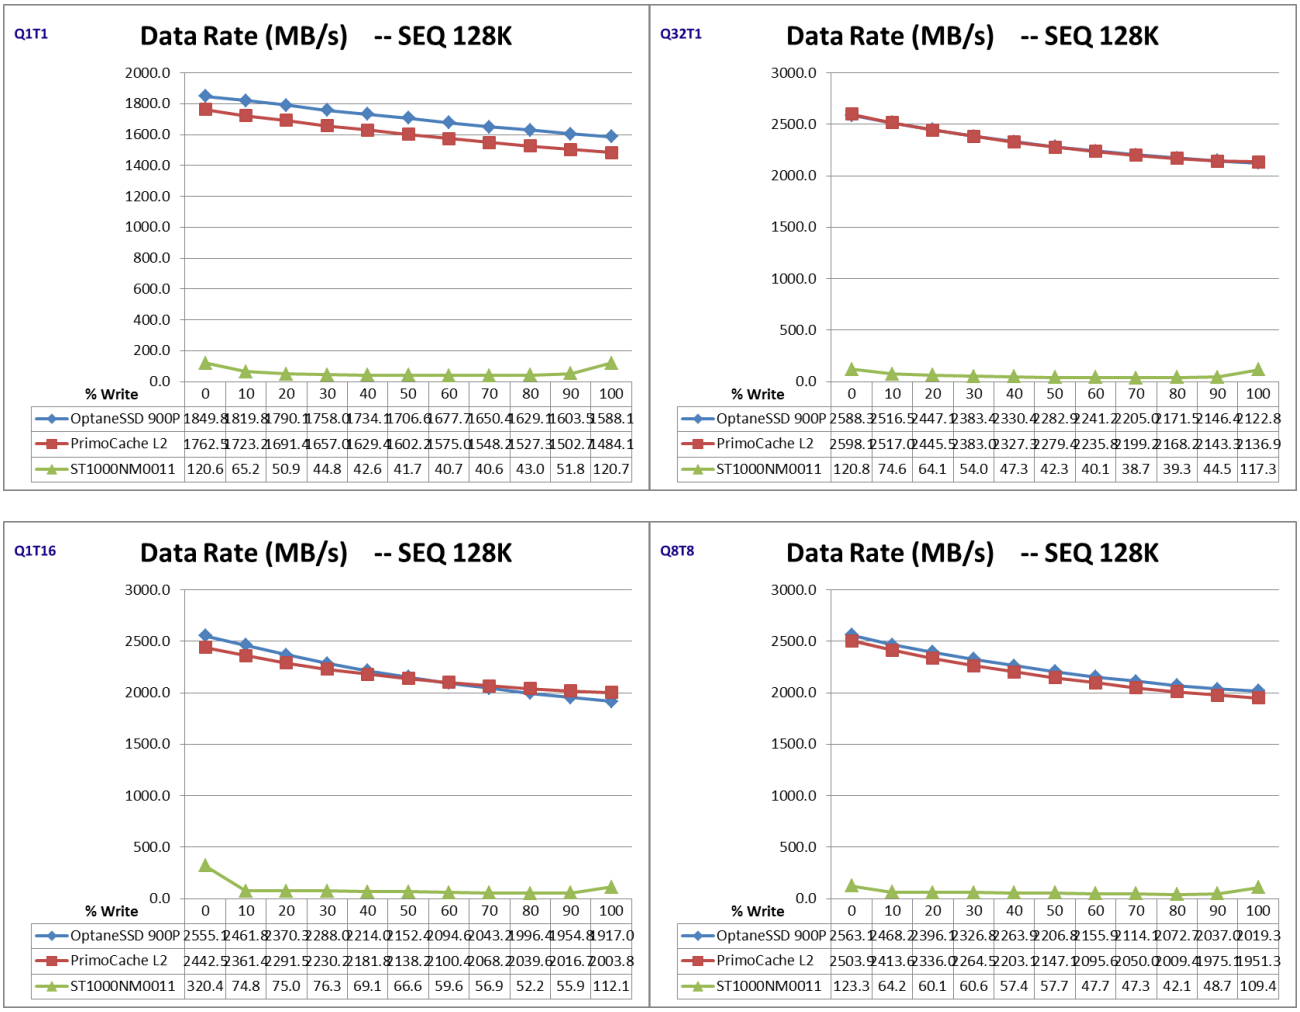

Q1T1	% Write	0	10	20	30	40	50	60	70	80	90	100
	OptaneSSD 900P	1849.8	1819.8	1790.1	1758.0	1734.1	1706.6	1677.7	1650.4	1629.1	1603.5	1588.1
	PrimoCache L2	1762.9	1723.2	1691.4	1657.0	1629.4	1602.2	1575.0	1548.2	1527.3	1502.7	1484.1
	ST1000NM0011	120.6	65.2	50.9	44.8	42.6	41.7	40.7	40.6	43.0	51.8	120.7

Q32T1	% Write	0	10	20	30	40	50	60	70	80	90	100
	OptaneSSD 900P	2588.3	2516.5	2447.1	2383.4	2330.4	2282.9	2241.2	2205.0	2171.5	2146.4	2122.8
	PrimoCache L2	2598.2	2517.0	2445.5	2383.0	2327.3	2279.4	2235.8	2199.2	2168.2	2143.3	2136.9
	ST1000NM0011	120.8	74.6	64.1	54.0	47.3	42.3	40.1	38.7	39.3	44.5	117.3

Q1T16	% Write	0	10	20	30	40	50	60	70	80	90	100
	OptaneSSD 900P	2555.2	2461.8	2370.3	2288.0	2214.0	2152.4	2094.6	2043.2	1996.4	1954.8	1917.0
	PrimoCache L2	2442.5	2361.4	2291.5	2230.2	2181.8	2138.2	2100.4	2068.2	2039.6	2016.7	2003.8
	ST1000NM0011	320.4	74.8	75.0	76.3	69.1	66.6	59.6	56.9	52.2	55.9	112.1

Q8T8	% Write	0	10	20	30	40	50	60	70	80	90	100
	OptaneSSD 900P	2563.3	2468.2	2396.1	2326.8	2263.9	2206.8	2155.9	2114.1	2072.7	2037.0	2019.3
	PrimoCache L2	2503.9	2413.6	2336.0	2264.5	2203.1	2147.1	2095.6	2050.0	2009.4	1975.1	1951.3
	ST1000NM0011	123.3	64.2	60.1	60.6	57.4	57.7	47.7	47.3	42.1	48.7	109.4

Random Mixed Read-Write

The following charts show the test results for random mixed read-write of the 4KB block under Q1T1, Q32T1, Q1T16, and Q8T8. It can be seen that as the read/write ratio changes, the SSD 900P basically exhibits a linear change at low thread count, but exhibits a curve change at high thread count where the mixed read-write IOPS value is lower than its corresponding linear value, but the deviation is not too big.

The performance of the PrimoCache cache in 4KB random mixed read/write is similar to the test results of the 4KB standard read/write, which is no longer repeated here.

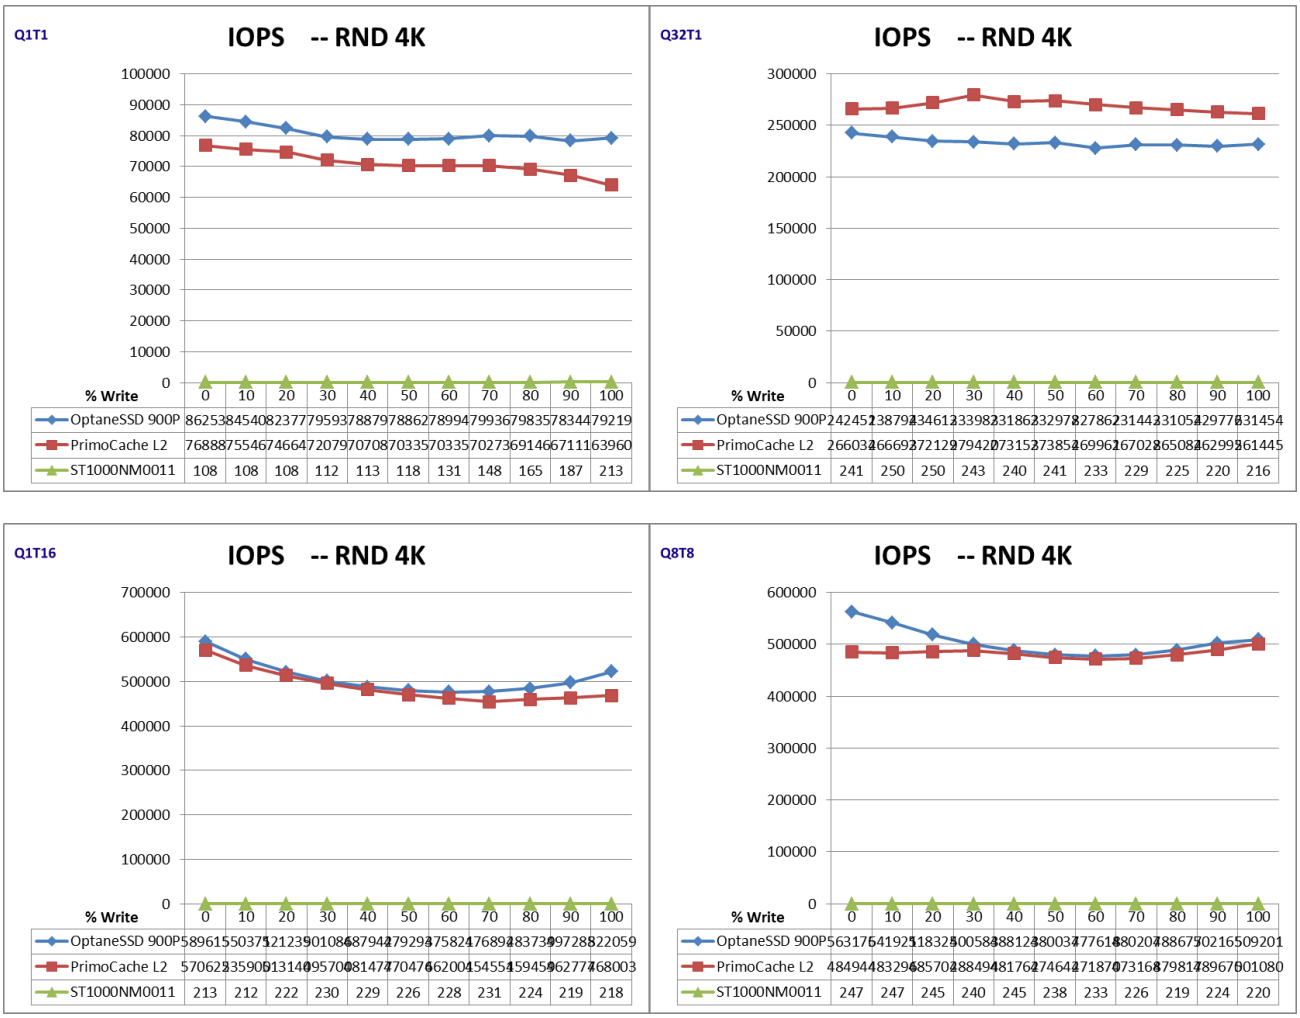

COMPLETE/PARTIAL SPACE TEST

In order to evaluate the performance difference of a storage device within different physical storage spaces, standard read/write tests are performed and compared for the entire storage space and the partial storage space of 10GB span in front and rear area of the Optane SSD 900P. All tests are performed under the condition that the entire SSD is filled with test data. “Full Drive” in the contrast charts of test results refers to the result of complete storage space, “Front 10G” refers to the 10GB space span in the front area, where the actual test selects the physical address space of 13GB - 23GB, and “Rear 10G” refers to the 10GB space in the rear area, where the actual test selects the physical address space of 223GB - 233GB.

The test results of different thread and queue depth are similar, so only some typical charts are selected here for analysis. From these charts, it can be seen that the performance of sequential read has very good consistency across the whole storage space, and the performance of sequential write is slightly different in different storage spaces. The 4KB random read/write IOPS also differ slightly with the address space, but the random read/write performance of the 8KB block or other sizes is basically the same. On the whole, the spatial consistency of the SSD 900P is quite good.

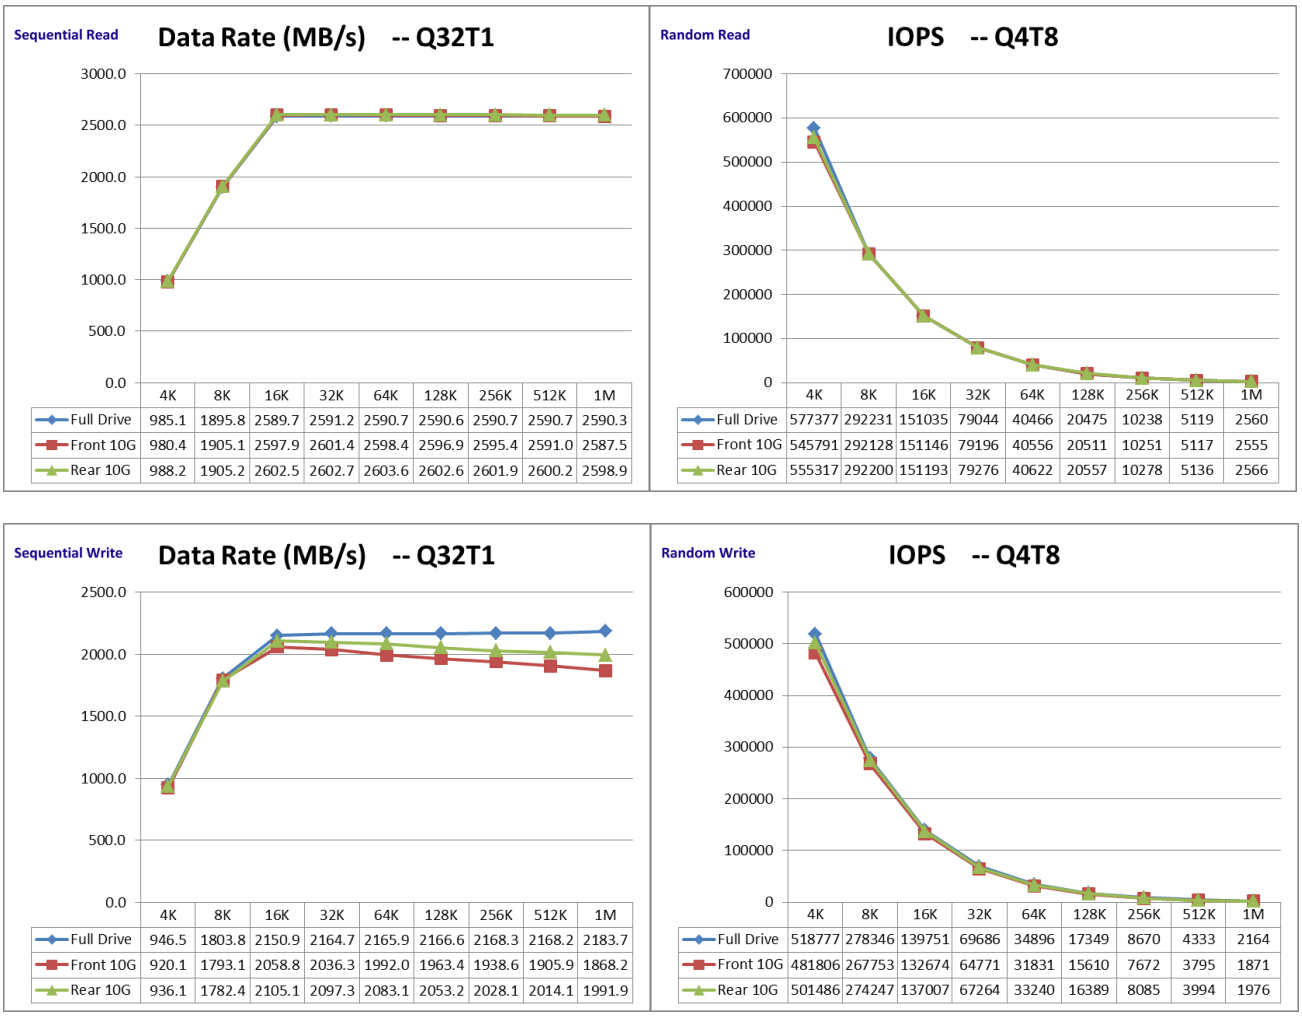

Sequential Read	4K	8K	16K	32K	64K	128K	256K	512K	1M
Full Drive	985.1	1895.8	2589.7	2591.2	2590.7	2590.6	2590.7	2590.7	2590.3
Front 10G	980.4	1905.1	2597.9	2601.4	2598.4	2596.9	2595.4	2591.0	2587.5
Rear 10G	988.2	1905.2	2602.5	2602.7	2603.6	2602.6	2601.9	2600.2	2598.9

Random Read	4K	8K	16K	32K	64K	128K	256K	512K	1M
Full Drive	577377	292231	151035	79044	40466	20475	10238	5119	2560
Front 10G	545791	292128	151146	79196	40556	20511	10251	5117	2555
Rear 10G	555317	292200	151193	79276	40622	20557	10278	5136	2566

Sequential Write	4K	8K	16K	32K	64K	128K	256K	512K	1M
Full Drive	946.5	1803.8	2150.9	2164.7	2165.9	2166.6	2168.3	2168.2	2183.7
Front 10G	920.1	1793.1	2058.8	2036.3	1992.0	1963.4	1938.6	1905.9	1868.2
Rear 10G	936.1	1782.4	2105.1	2097.3	2083.1	2053.2	2028.1	2014.1	1991.9

Random Write	4K	8K	16K	32K	64K	128K	256K	512K	1M
Full Drive	518777	278346	139751	69686	34896	17349	8670	4333	2164
Front 10G	481806	267753	132674	64771	31831	15610	7672	3795	1871
Rear 10G	501486	274247	137007	67264	33240	16389	8085	3994	1976

WRITE-LIFE IMPACT TEST

Because SSDs are usually limited in number of write cycles, to test the performance changes after a large number of writes, this article introduces the write-life impact test on the SSD 900P, evaluating the impact of SSD write life on performance. Due to the test on the entire drive is very time consuming, which requires more than 30 days to reach the claimed lifetime of the SSD 900P (280GB) even if the drive is uninterrupted written at the maximum speed, only 10 GB of storage space is selected for the actual test. The claimed endurance-per-GB of SSD 900P is 18.69TBW, so the endurance of 10GB is 186.9TBW. In this space, the standard performance test was performed after the uniform write of 0TB (0% Worn), 90TB (50% Worn), 180TB (100% Worn) and 200TB (110% Worn).

From the test result comparison chart, it can be seen that the sequential read speed, taking Q32T1 as an example, decreases from 2,600MB/s to 2,350MB/s after 50% worn, which only falls less than 10%. The performance is finally stable at 2,270MB/s, only less than 13% down, when it reaches 100% of lifetime or above. The 4KB random read IOPS performance, taking Q16T16 as an example, remains unchanged even at 110% worn. This is a very surprising test result. The sequential write performance (Q32T1) and the 4KB random write IOPS (Q16T16) have performance reductions of only about 10% and 13% respectively after 50% worn. What's odd is the write performance at 100% worn or above is a little better than that at 50% worn.

After reaching 110% worn, we verified the accuracy of read/write data for this storage space. After many tests, the data that was written and then power-off to read back was consistent with the source data, and no errors occurred. According to the test results, it seems that the Intel claimed endurance value of the SSD 900P is kind of conservative, and **the actual life expectancy should exceed the claimed value**.

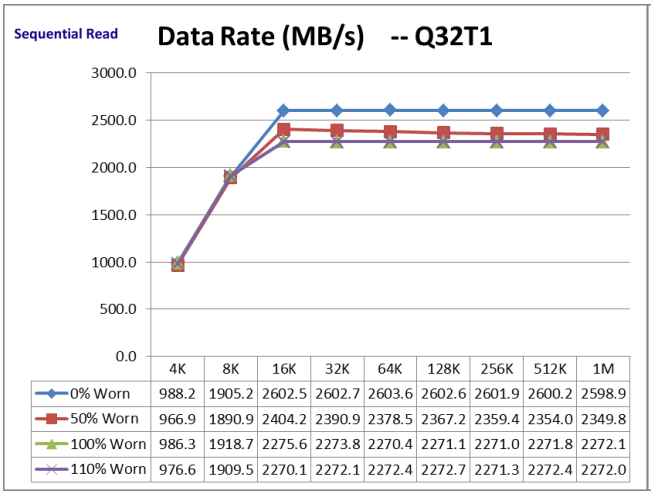

	4K	8K	16K	32K	64K	128K	256K	512K	1M
0% Worn	988.2	1905.2	2602.5	2602.7	2603.6	2602.6	2601.9	2600.2	2598.9
50% Worn	966.9	1890.9	2404.2	2390.9	2378.5	2367.2	2359.4	2354.0	2349.8
100% Worn	986.3	1918.7	2275.6	2273.8	2270.4	2271.1	2271.0	2271.8	2272.1
110% Worn	976.6	1909.5	2270.1	2272.1	2272.4	2272.7	2271.3	2272.4	2272.0

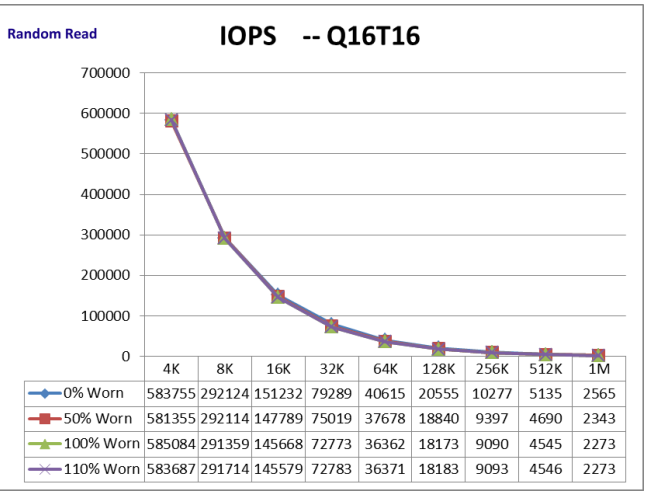

	4K	8K	16K	32K	64K	128K	256K	512K	1M
0% Worn	583755	292124	151232	79289	40615	20555	10277	5135	2565
50% Worn	581355	292114	147789	75019	37678	18840	9397	4690	2343
100% Worn	585084	291359	145668	72773	36362	18173	9090	4545	2273
110% Worn	583687	291714	145579	72783	36371	18183	9093	4546	2273

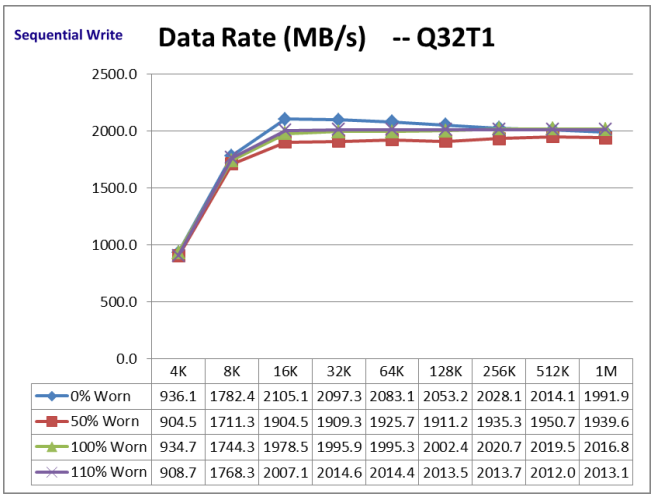

	4K	8K	16K	32K	64K	128K	256K	512K	1M
0% Worn	936.1	1782.4	2105.1	2097.3	2083.1	2053.2	2028.1	2014.1	1991.9
50% Worn	904.5	1711.3	1904.5	1909.3	1925.7	1911.2	1935.3	1950.7	1939.6
100% Worn	934.7	1744.3	1978.5	1995.9	1995.3	2002.4	2020.7	2019.5	2016.8
110% Worn	908.7	1768.3	2007.1	2014.6	2014.4	2013.5	2013.7	2012.0	2013.1

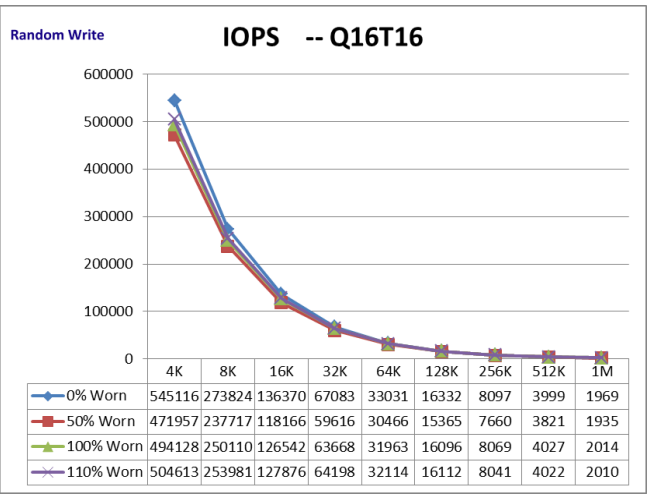

	4K	8K	16K	32K	64K	128K	256K	512K	1M
0% Worn	545116	273824	136370	67083	33031	16332	8097	3999	1969
50% Worn	471957	237717	118166	59616	30466	15365	7660	3821	1935
100% Worn	494128	250110	126542	63668	31963	16096	8069	4027	2014
110% Worn	504613	253981	127876	64198	32114	16112	8041	4022	2010

CONCLUSIONS

Combined with the above test results, it can be seen that the measured results of various performance items of the Optane SSD 900P can reach the official claimed value, and even most of the results exceed the official value. Its performance on read/write speed, IO throughput, and access latency is very strong, far beyond the mainstream SSDs in current market. Especially under the medium/heavy workload, its performance can be fully utilized.

The Optane SSD 900P maintains performance consistency across mixed read-write tests, different physical storage space, and at different lifetime stages, without significant performance degradation. The measured performance is only degraded about 10%, no more than 15%, after the official write lifetime is reached.

It can also be seen that when the SSD 900P is served with the PrimoCache software as a cache for other drives, the performance of cached drives is generally close to or reach the performance of the SSD 900P. In some cases, it even exceeds the performance of the SSD 900P.

Combining the excellent read/write performance and high stability of the SSD 900P and the cache performance of the PrimoCache software, in addition to applications in the desktop consumer market, the SSD 900P is also suitable as a cache for those large-capacity slow hard drives on a workstation or server, which is easy and convenient to improve the read/write speed and IO throughput of hard drives, to make it close to or reach, even exceed the performance of the SSD 900P. It is ideal for users who have budget restraints or do not want to modify the system configuration or to do the data migration.

Although this article does not test the practical performance improvement of specific applications of the PrimoCache cache on the server, the test scheme here covers the various IO processing scenarios that a storage device may encounter in specific applications, and shows the benchmark performance of the storage device. Taking SQL Server as an example, the following table which comes from Diskspd help documentation lists all kinds of possible SQL Server operations. Therefore, to evaluate the drive IO performance on a SQL Server system, a quantitative result can be obtained by simply measuring the corresponding random or sequential access performance of related data block patterns.

File Type	Operation	READ pattern	WRITE pattern	Threads used	I/O type
Data File	Normal Activity	8KB - 128KB	8KB - 128KB	Based on MaxDOP	Random
	Checkpoint	N/A	64KB - 128KB	# of Sockets in Computer	Random
	LazyWriter	N/A	64KB - 128KB	1 per NUMA Node	Random
	Bulk Insert	N/A	8KB - 128KB	Based on MaxDOP	Sequential
	Backup	1 MB	1 MB	Based on MaxDOP	Sequential
	Restore	64KB	64KB	Based on MaxDOP	Sequential
	DBCC Checkdb w/ no repair option	8KB - 64KB	N/A	Based on MaxDOP	Sequential
	Rebuild Index	Up to 512 KB	8KB - 128KB	Based on MaxDOP	Sequential
	ReadAhead	Up to 512 KB	N/A	Based on MaxDOP	Sequential
Log File	Normal Activity	512B - 64KB	512B - 64KB	one log writer thread per soft NUMA node with a cap of 4	Sequential

It is worth mentioning that PrimoCache also supports RAM as a cache for hard drives at the same time. In order to accurately reflect the performance of using the Optane SSD 900P as cache, this article tested the PrimoCache cache performance without setting up RAM cache, but instead used the SSD 900P purely. In practical applications, RAM can be used as a level-1 cache, in addition to the SSD 900P as a level-2 cache, for better performance.

Of course, in order to test the maximum performance of the PrimoCache cache, the tests in this article were performed under the premise that the test files are completely cached to the SSD 900P. In practical applications, because the cache space is usually limited, it is not guaranteed that all data or files in cached drives will be cached. In addition, the cache configuration may vary according to actual requirements. So it is possible that the performance in practical applications and test results will be different from this article.

APPLICATION SCENARIOS

From the above test results and conclusions, by the solution of using PrimoCache software and Optane SSD 900P as a high-speed and large-capacity cache for slow storage devices in the system, both the read and the write performance of the system can be greatly improved. So it is not difficult to find that this solution is able to be adopted in a very broad application scenarios and fields. It can easily solve the problem of read/write performance bottlenecks in workstations, small and medium servers. If you want to build a new system with high access performance and high storage capacity, then this solution plus large-capacity slow drives offers a superior price-performance ratio compared to the solution that directly uses large-capacity high-speed drives. The following is a simple list of working scenarios to which this solution can be applied, but not limited to these scenarios.

- Workstations: For example, for a graphics workstation that is mainly used to render animations from 3DMAX, MAYA, C4D and other 3D software, after applying the SSD 900P caching solution, the I/O efficiency of the entire system can be greatly improved, thus the rendering time can also be greatly reduced.

- Rendering Farms: During the rendering process, render nodes will generate a large number of near-simultaneous I/O requests to the server. The I/O performance of the conventional storage system is difficult to meet such demand, especially in the case of a large number of render nodes. However, the SSD 900P cache solution will solve this problem very well.

- VDI Systems: When a large number of thin clients or diskless stations boot up at the same time, a large number of read operations will be performed simultaneously on the server drives. With the characteristics of high throughput, low latency and long service life of the SSD 900P, combined with PrimoCache software, the caching solution can reduce the cost of the whole VDI system.

- Traditional NAS or Distributed Storage Solutions: The SSD 900P caching solution can be adopted to accelerate storage performance and upgrade the system simply and safely.

- Diskless Internet Cafes: For example, the current popular game *PlayerUnknown's Battlegrounds* in Internet cafes will generate a large amount of data to write back to the server. Traditional SSDs will reach the limit of write cycles soon due to heavy write workload. The SSD 900P with the features of ultra-high endurance and ultra-fast write, as a write-back drive cache for the diskless Internet cafe servers, can effectively solve this problem!

APPENDIX: TEST REPORTS

STANDARD READ/WRITE TEST REPORTS

Sequential Read

Sequential Write

Random Read

Random Write

MIXED READ-WRITE TEST REPORTS

Q1T1

Q32T1

Q1T16

Q8T8

COMPLETE/PARTIAL SPACE TEST REPORTS

Sequential Read

Sequential Write

Random Read

Random Write

WRITE-LIFE IMPACT TEST REPORTS

Sequential Read

Sequential Write

Random Read

Random Write