All Flash Virtual SAN Solution

MEMBLAZE

Boost User Experience While running Oracle DBs

- 10K TPS for Oracle TPC-C Test per node
- 322MBPS random write per node
- 30K TPS on 3 nodes with ability to scale to 64 Virtual SAN nodes
- Average IO latency of PBlaze4 is under 0.4 ms

Optimize Return on Investment

- Combined computing and storage， Reduces storage costs significantly
- Distributed architecture can expand horizontally and flexibly and deploy quickly according to the business requirements.
- Increase productivity by reusing vSphere management and storage services

Industry Leading Components used for All-Flash Virtual SAN

- 3*Dell's PowerEdge 13G Servers(2*Intel Xeon E5-2630 v3 2.4GHz,16Cores & 24 Threads, 128 GB, 10 Gigabit Network Card Intel 82599EB , Memblaze PBlaze4 PCIe SSD 800GB, Memblaze PBlaze4 PCIe SSD 3.2TB)
- 1*HuaWei S5700(One Gigabit switch with four 10 Gigabit ports)

The Architecture of Virtual SAN Solution Design

The construction characteristics of Virtual SAN:

- The base configuration of Virtual SAN is 3 or 2 nodes with Witness components;
- Each node can provide computing and storage capacities, but also capacity of computing only or storage only;
- Flash memory plays a pivotal role in the performance of Virtual SAN architecture, which possesses the characteristics of high IOPS, high bandwidth and low latency. It can greatly optimize the performance of Virtual SAN system with the Virtual SAN hotspot capturing ability, which can meet the increasingly complexity of business virtualization storage needs;

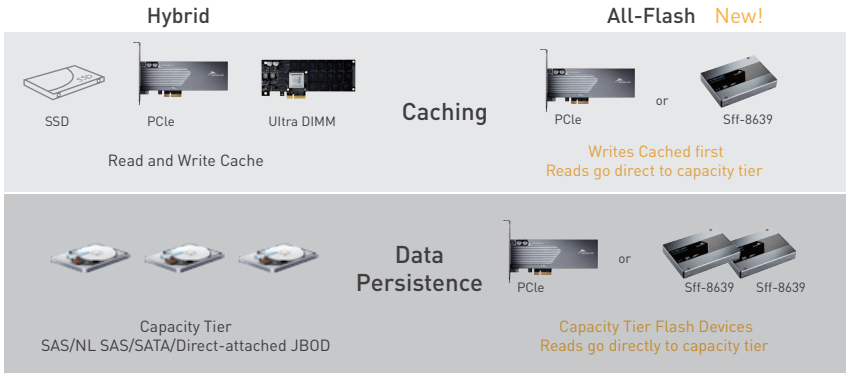

Virtual SAN Architecture

Virtual SAN is a high-level integrated architecture Server Virtual SAN, with two different architectures:

Hybrid Architecture

The hybrid architecture is characterized by high performance cost ratio and high system capacity. Under the hybrid architecture, flash SSD is used as a read and write cache with an obvious optimizing for the hotspot data to be read, and can carry out the unexpected write operation with the flash memory layer to meet the business requirements (such as the boot storm and the antivirus storm under VDI environment).

According to the accumulated experience of hybrid architecture Virtual SAN using by the clients of Memblaze, the sum of each node's SSD capacity should be larger than the hot data of client business system in Virtual SAN, which can ensure the IO processing of hot data is completed by flash memory.

All Flash Memory Architecture

The all flash memory architecture after Virtual SAN 6.0 can use the SSD in the cache layer and capacity layer, which means the clients can take one SSD as the cache layer of Virtual SAN and another SSD of the same type as the capacity layer of Virtual SAN.

The all flash memory architecture is featured by persistent read and write, and can guarantee the high performance. No matter the IO in cache layer or capacity layer, the response time of IO is flash level.

The architecture with two layers of flash memory can improve the speed of data moves from cache layer to capacity layer. The impact on the performance of the front-end application is smaller than the hybrid architecture's impact when the data moves from cache layer to capacity layer.

Compared with hybrid architecture, the performance of the capacity layer in all flash memory architecture is very high, which almost without the need of striping of crossing nodes and means lower network latency appears.

Virtual SAN+PBlaze4 All-Flash Solution Design

- Cache: 800GB in both 2.5" and Add-in Card (AIC) form factors.
- Capacity Layer : 1.6 to 3.2TB in both 2.5" and AIC or 6.4TB AIC
- Drive Type: NVMe

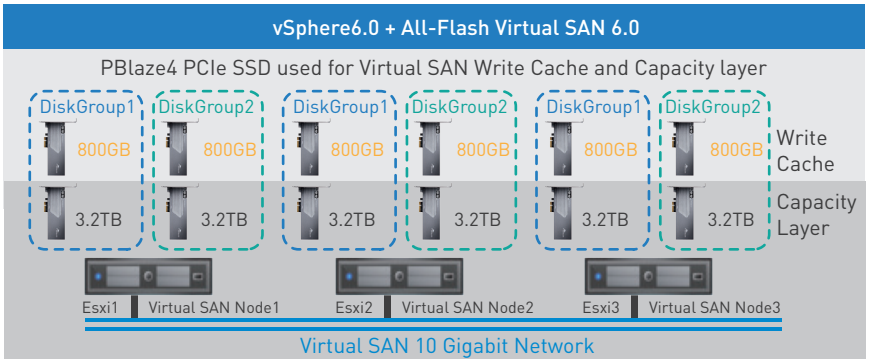

The diagram above shows the system architecture of all-flash Virtual SAN. the performance of the each diskgroup to the VMs as following table described:

3 Virtual SAN nodes with 6 disk groups Total Working Size=4.8TB *2(FTT=1)			
	Total VMs' Performance	Virtual SAN Performance	Performance per disk group
8KB Random Read	312.4k IOPS	312.4k IOPS	52.1k IOPS
8KB Random Write	117.6k IOPS	235.3k IOPS	39.2k IOPS
8KB Random R/W(7:3)	151.7k IOPS	197.2k IOPS	32.9k IOPS
64KB Sequential Read	4.32GB/s	4.32 GB/s	720 MB/s
64KB Sequential Write	1.84GB/s	3.68 GB/s	614 MB/s

PBlaze4+Virtual SAN6.0(All Flash) +Oracle VM Practice

- 3 Oracle HammerDB VM, 3*2000 WareHouse;
- The Volume of Valid Test Data: 2.79TB;
- Actual Usage of Virtual SAN(FTT=1): 5.45TB;
- 3. 32GB SGA, 3*135 concurrent user connections
- The result of TPC-C test: 3.74M TPM

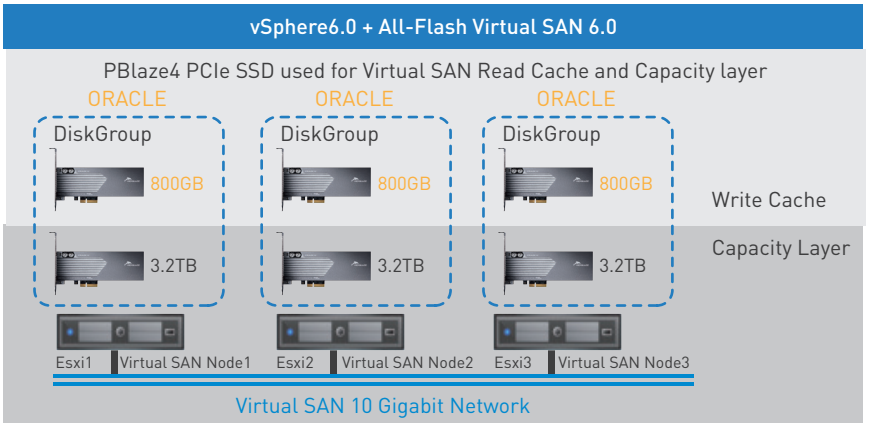

According to the AWR reports of the three Oracle VM, the HammerDB TPC-C test result can reach 31k TPS on the Virtual SAN environment. During the test, Virtual SAN storage policy is 'FTT=1', write IO pressure doubled, so there are 353 MBPS(44,756 IOPS) read and 965 MBPS(50,259 IOPS) write IO on the whole Virtual SAN system.

Test Items (the number of virtual users is 405)	Total metrics of the three Oracle VM (per second)
Redo size:	148.7MB
Logical read(blocks):	378k
Block changes:	100k
Physical read(blocks):	17k
Physical write(blocks):	23k
Transactions:	31k

Beijing Memblaze Technology Co., Ltd.
Address: 3F,Buliding B2, Dongsheng Park,66 Xixiaokou Road,Haidian,District,Beijing.
Hotline: 400-650-7995 contact@memblaze.com