



Scaling In-Memory Data Processing with Samsung Advanced DRAM and NAND/SSD Solutions

Samsung and VoltDB analyze advances in memory and in-memory database technologies to address challenges from massive increases in contextual data

A White Paper by Samsung Memory Solutions Lab and VoltDB, Inc.

Scaling In-Memory Data Processing with Samsung Advanced DRAM and NAND/SSD Solutions

Executive Summary

Rapidly evolving data utilization patterns, in accommodating greater volumes of contextual as well as rich streaming data, are leading to the use of larger datasets for faster, transactional workloads. Seeing this early in the evolutionary cycle, Samsung has been leading the effort to create high-performance, high capacity memory for servers that will help to decrease datacenter operational expenses. These breakthroughs include high-density RAM, and faster storage via Samsung NVMe technology. VoltDB, an in-memory operational database that optimizes decision-making from fast data, benefits from the use of high-density RAM for system main memory, and NVMe technology for durability.

This white paper describes work conducted by Samsung and VoltDB, Inc. which demonstrates that high-capacity DRAM can be leveraged by the VoltDB database to deliver high-throughput with large in-memory datasets. The improved performance delivered by the combined use of VoltDB and Samsung memory scales with the capacity of the system. Samsung NVMe SSDs provide a high-performance durability layer that allows for snapshot and recovery of any node in the system. The reader will be introduced to use cases and benchmarks established by Samsung that demonstrate why Samsung NVMe SSDs represent a superior solution for deploying large-scale, in-memory storage at this layer, while enabling seamless back-ups.

Introduction

In-memory computing is here to stay. After a decade of development, it has come of age at a most opportune time. In-memory has evolved from solutions that offered low throughput and low latency, to systems that offer high throughput and impressively lower latency, and now to systems with exceptionally high-throughput reads and even lower latencies.

Three fundamental factors have been driving this trend:

- Hardware is now available that can use higher amounts of DRAM (High Capacity DRAM) – up to 3 TB – as compared to the previously existing barrier of around 512 GB RAM per server
- SSDs have become a ubiquitous solution for all datacenter storage needs. While SATA SSDs cannot meet performance SLOs (service level objectives) at the higher end of the application spectrum, this can be accomplished with advanced NVMe storage technology and Samsung's recently developed Z-SSD (5X faster than NVMe).
- The benefits of very large servers with high-density DRAM and high-speed SSDs can be fully attained when systems are designed to take advantage of them.

VoltDB: An In-Memory, Real-Time Decision-Support System

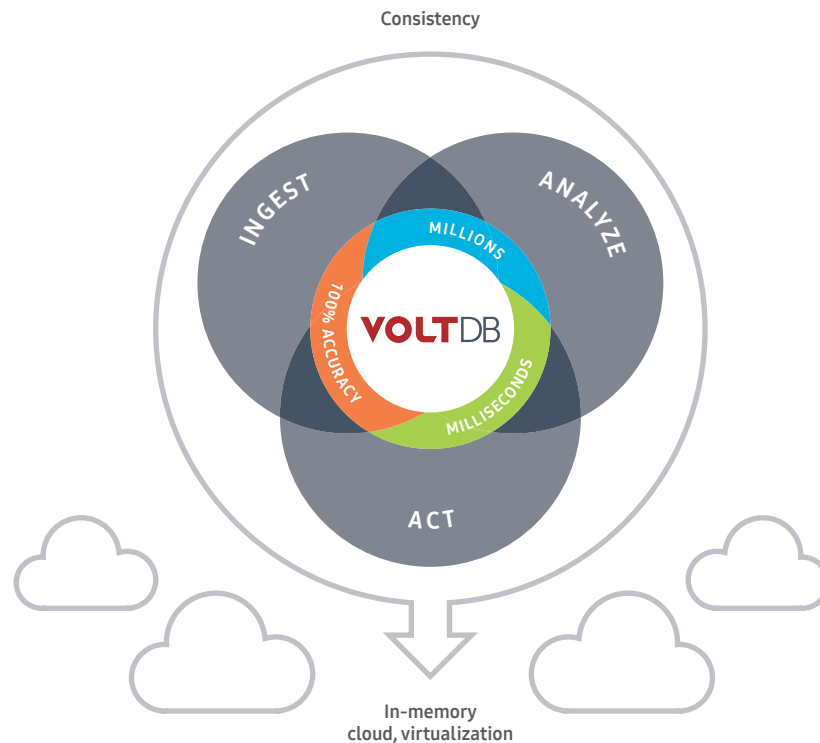


Figure 1 – Confluence of Fast Data requirements met by a single layer - VoltDB.

VoltDB

VoltDB has advanced the role of an OLTP database into that of a real-time DSS (Decision Support System). It has already enabled this transformation in a variety of use cases. These include cases involving:

- real-time portfolio risk monitoring and management
- credit card fraud detection and prevention
- monitoring of sensor data for anomaly detection
- monitoring of network management data
- several elements of the Evolved Packet Core (mobility management – MME, user authorization and authentication – AAA, policy management – PCRF/PCEF, Least Cost Route calculations, real-time billing – OCS/OFCS)
- other key functions of telecommunications companies
- hyper-personalized customer interaction applications typically found in the online gaming and advertising industries.

There is a new class of applications and use cases that requires ultra-fast interaction with the database, within the context of a

larger dataset than would be traditionally considered “memory-friendly.” These applications require single-digit millisecond latency across data sets in the 10s of terabytes.

Today’s conventional scale up vs. scale out data processing is now evolving into a postmodern definition of scale up AND scale out. Change agents pushing this transformation range from Single Source of Truth (combining multiple ODSs) initiatives in traditional financial service enterprises, as well as the increasing desire of ecommerce organizations to require extremely fast data collection, and to significantly improve next-generation filtration systems for the Industrial Internet of Things (IIoT).

Samsung has been spearheading a rapid evolution in chip architectures (including those pertaining to High Capacity DDR4, NVMe SSD and NAND/Z-SSD) that is enabling a critical advancement in the design of in-memory data processing systems like VoltDB. These advanced memory developments are able to scale to handle ever more complex applications with larger datasets and greater volumes of contextual data, without compromising data safety or recovery time SLOs).

Samsung Memory Technologies

High Capacity DRAM

Samsung High Capacity memory solutions are well exemplified by today's most advanced DDR4 technology. Modules comprised of Samsung DDR4 are designed with new system circuit architecture to deliver significantly higher performance with lower power requirements than previous generation of memory products.

Doubled bandwidth, along with reduced voltage and more efficient power consumption, not only improves performance, but also optimizes total cost of ownership. Furthermore, the deployment of Samsung 8Gb DDR4, produced using an advanced 10nm-class process, assures an optimal level of performance, greater bandwidth, sizeable power savings and higher reliability in the manufacturing of high capacity DIMMs.

Today, Samsung DDR4 DRAM offerings range from 8GB capacity points to 128GB capacity points (Table A). Samsung DDR4 LRDIMM (load reduced memory module) technology uses a distributed

buffer to maximize memory bandwidth efficiency when scaling to higher capacities and speeds. In general, RDIMMs are more power-efficient, while Load Reduced DIMMs show an improvement in speed.

Capacity	DIMM
8GB	RDIMM
16GB	RDIMM
32GB	RDIMM/LRDIMM
64GB	RDIMM(*)/LRDIMM
128GB	RDIMM(*)/LRDIMM(*)

Table A. Samsung DDR4 Capacity Points

As a result, the combination of LRDIMM's features and Samsung 8Gb DDR4 chips produced using a 10nm process node provides an optimized solution for highly virtualized environments, high-performance computing and networking.

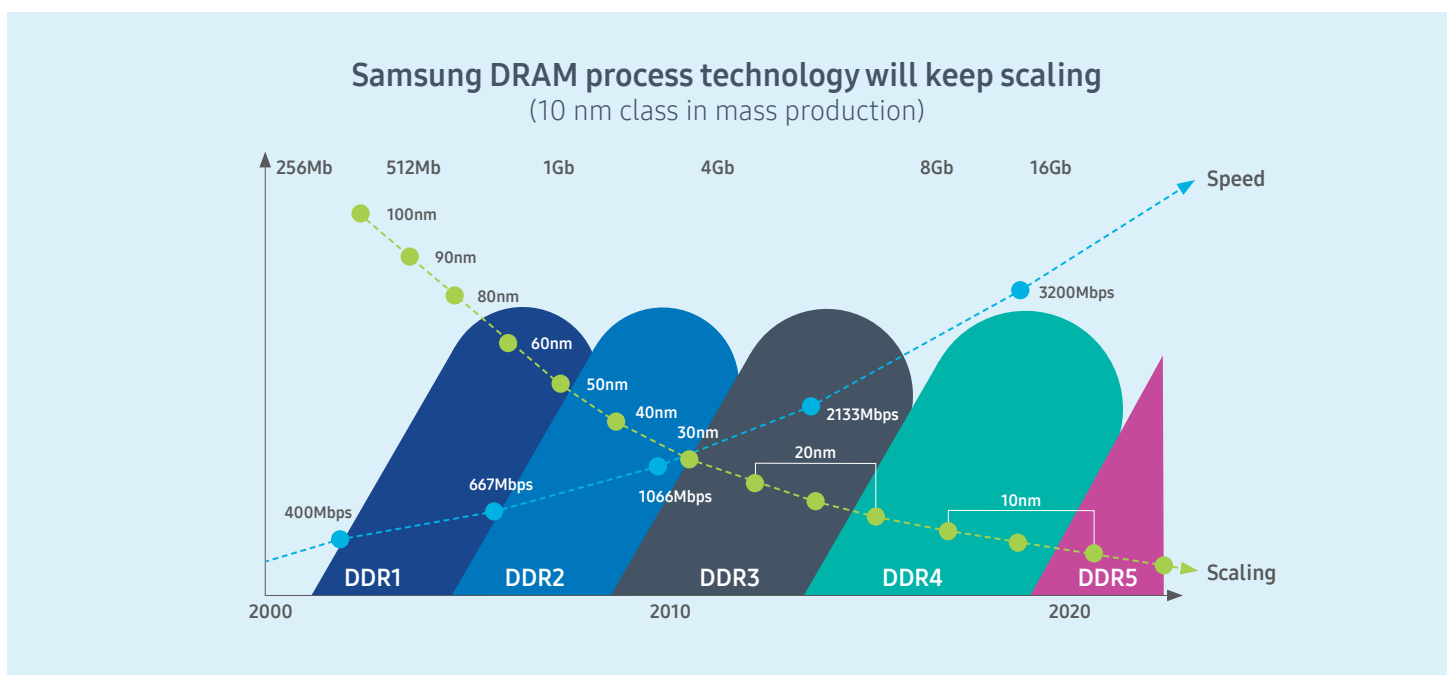


Figure 2. Samsung DRAM Process Technology

In this study, we have leveraged a Samsung 64GB, 8Gb-based DDR4 LRDIMM memory solution operating at 2133MHz to provide an attractive high performance alternative for those seeking a more cost-efficient solution.

Samsung Memory Technologies (continued)

PM1725 NVMe SSD

To satisfy the exceptionally high demands of enterprise environments, SSDs must perform over long periods at their maximum levels to accommodate the variety of workloads simultaneously accessing the devices. In addition, the selected SSD must provide consistent performance that will satisfy highly demanding Quality of Service (QoS) requirements. The Samsung PM1725 SSD, which was selected for these tests, is designed to excel in virtually any datacenter scenario. This enterprise-level, ultra-high-performance SSD provides unsurpassed random read performance and is particularly well suited for read-intensive datacenter applications. When compared with the other standardized SSDs in Samsung internal testing, the PM1725 delivers the highest random read IOPS.

INCREASED PERFORMANCE



		PM863A	PM963	PM1725	
	Interface	SATA	NVMe		
	Form Factor	2.5"	2.5"	M.2	2.5" HHHL
	Capacity (Max)	3.84TB			6.4TB
	DWPD	1.3 for 3 Years			5 for 5 Years
Performance	Sequential R/W (Max) MB/s	520/480	2K/1.2K	3.1K/2K	6K/2K
	Random R/W (Max) IOPs	97K/24K	430K/40K	750K/120K	1000K/120K

Figure 3. SSD performance



PM1725a / PM1725 is embedded with cutting-edge flash memory leveraging V-NAND technology. PM1725a / PM1725 uses a state-of-the-art NVMe interface to overcome all the bottlenecks experienced in SAS and SATA interfaces. Furthermore, PM1725a / PM1725 empowers the highest density to successfully process intensive and heavy enterprise workloads.

VoltDB & Samsung Benchmarks

VoltDB partnered with Samsung to tune its software so that it could take full advantage of Samsung's chip architecture and speeds. In exploiting the performance benefits of this combination, we conducted benchmarks that leveraged the attributes of the Samsung hardware. Performance metrics of particular interest were the following:

1. Transactions per second and Transactions per minute, where appropriate
2. 95 percentile response latency
3. Mean time to recover from failures

Single Node Configuration

To demonstrate the benefits of running larger-than-memory datasets, we ran benchmarks against two workloads, Voter and TPC-C.

Model	PowerEdge R730xd
OS	Linux 4.20-42 generic — Ubuntu release 15.10 wily
CPU	E5-2699 v4 @2.20GHz
Processors	22
Sockets	2
Threads	88
Total DRAM	1TB
Storage (NVMe)	PM1725 (1.6TB)
Storage (SATA)	SM863a (1.8TB)
Network	Mellanox SN2700 (25Gbe)

Table B. Single Node Configuration

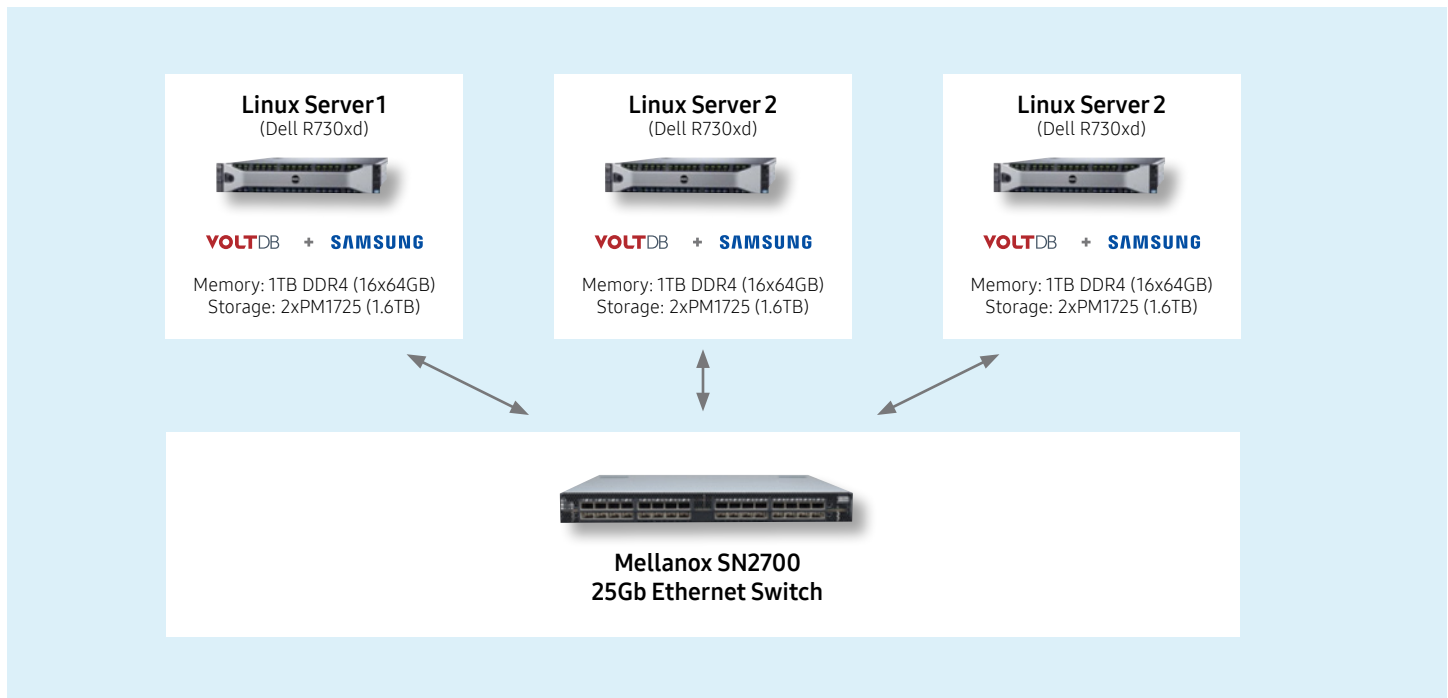


Figure 4. Cluster Configuration

VoltDB & Samsung Benchmarks (continued)

Voter and Voterw/Survey

This example simulates a phone-based election process. Voters (based on phone numbers generated randomly by the client application) are each allowed a limited number of votes. When voting with Survey, the sample application uses the same set of transactions as VOTER, with the exception of an additional column having been added to the table "VOTES." The extra column stores answers and comments to the survey that callers take when they vote. These workloads represent use cases where a large amount of data is stored, aggregated and incorporated with incoming data to drive fast decision-making with low latency responses.

Other use cases could be:

1. Single source of truth with account snapshot recalculation in core-banking scenarios
2. Real-time pre-authorization requests in the processing of health care payers, where the entire history of patient information would be considered contextual
3. Advertisement targeting where many advertisers, tags and audiences need to go through a match-making process. Greater accuracy yields higher revenue.
4. Telco and IIoT applications that have to contend with a barrage of data needing real-time processing, filtration and aggregation.

TPC-C

VoltDB's implementation of the TPC-C benchmark is robust, and adheres to the standard closely. We tested initial loading and the full transaction set defined by the benchmark conducted here. This could be beneficial for any e-commerce organization looking to create a single consistent storage and management platform for its orders, inventory/catalog and fulfillment data.



Performance Results

Voter w/ Survey

The graphs below show that this workload achieves 2M transactions per second with a 2TB database across three nodes. However, here there is no durability of the database. VoltDB supports using both NVMe SSD and SATA SSD drives to provide an acceptable level of durability. When high durability is required, Figure 5 and Figure 6 show that Samsung NVMe SSD delivers almost 4X throughput improvement over SATA drives, while maintaining a low response time of 1/3X that of SATA drives.

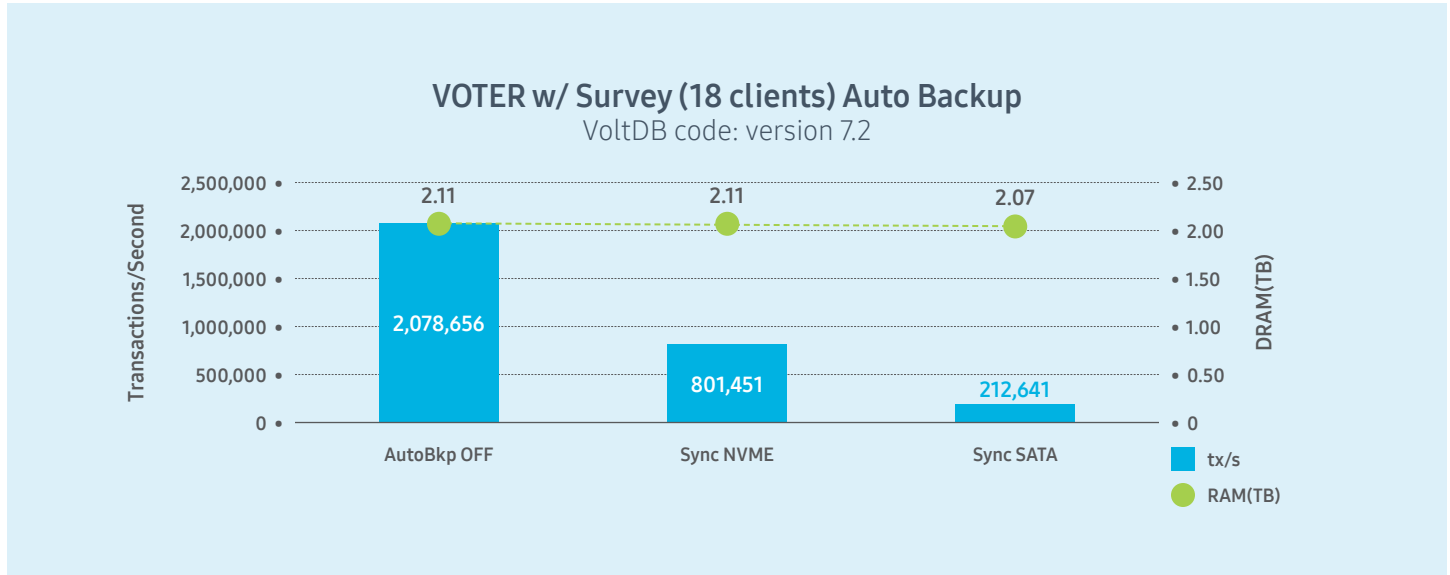


Figure 5. Voter w/ Survey throughput




Figure 6. Voter w/ Survey latency

Performance Results (continued)

TPC-C / On-line Transaction Processing Benchmark

VoltDB can deliver an incredible 5M New Order transactions per minute for a 1000-warehouse TPC-C database, but when we add the “Auto Backup” feature to provide durability, performance and latency degrade, as expected. However, our results show that NVMe SSDs deliver 80% more throughput at a 12ms response time compared to SATA SSDs.

VoltDB is an in-memory database that benefits from high capacity DRAM. If durability is required, as well as strict SLO, applications running on VoltDB will need fast storage, which is delivered by enterprise class NVMe SSDs from Samsung.




Figure 7. TPC-C throughput and latency

Conclusion

Today's rapid evolution of data utilization patterns is leading to greater amounts of data being used for faster processing, and even transactional work.

Samsung has been leading efforts to enable the creation of larger-footprint servers that will decrease datacenter operational expenses, with its high-density DRAM and ultra-fast NVMe SSD storage drives.

For more information contact Samsung at:
msl-inquiry@ssi.samsung.com

VoltDB, Inc. has been a pioneer in rapidly and accurately extracting value and decisions/actions from fast data. Its solutions have been enhanced with the use of high-density DRAM for the server main memory, and NVMe SSD technology for greater durability.

For more information contact VoltDB at:
info@voldb.com

Copyright 2017 Samsung Electronics, Co. Ltd. and VoltDB, Inc. All Rights Reserved. All brand, product, service names and logos are trademarks and/or registered trademarks of their respective owners and are hereby recognized and acknowledged. Specifications and designs are subject to change without notice. All data were deemed correct at time of creation. Samsung and VoltDB are not liable for errors or omissions. PID 2017-07