

The logo for the Storage Networking Industry Association (SNIA), consisting of a small square icon followed by the letters 'SNIA' in a bold, sans-serif font.

# PERSISTENT MEMORY PMM SUMMIT

JANUARY 18, 2017 | SAN JOSE, CA

## Bringing Persistent Memory Technology to SAP HANA: Opportunities and Challenges

Zora Caklovic, Product Expert and Oliver Rebholz, Development Manager  
SAP HANA Products & Innovation, SAP SE

# Disclaimer

This presentation outlines our general product direction and should not be relied on in making a purchase decision. This presentation is not subject to your license agreement or any other agreement with SAP. SAP has no obligation to pursue any course of business outlined in this presentation or to develop or release any functionality mentioned in this presentation. This presentation and SAP's strategy and possible future developments are subject to change and may be changed by SAP at any time for any reason without notice. This document is provided without a warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. SAP assumes no responsibility for errors or omissions in this document, except if such damages were caused by SAP intentionally or grossly negligent.

# Agenda

Persistent Memory (PM): What's all the noise about?

Bringing PM to in-memory computing: Opportunities & Challenges

Leveraging Intel DIMM based on 3D XPoint™ technology for SAP HANA

Sample use case: Improved in-memory database restart time

Conclusions

# Persistent Memory: What's all the noise about?

Persistent Memory (PM): An emerging, next-generation non-volatile memory technology

- A true game changer for in-memory computing offering new opportunities for increased performance and scalability and lowering cost


Key business drivers for bringing PM to in-memory computing:

- As data volumes keep growing, in-memory databases like SAP HANA require larger capacities
- DRAM scaling has slowed significantly and will offer lower capacities
  - New solutions needed for cost-efficient store of large data while still providing real-time access

## Persistent Memory

- Byte addressable persistence
  - Fast enough to load directly
  - Usually on memory bus
- NVDIMMs available today
- 3D XPoint™ Memory
  - Persistent
  - (up to) 1000X faster than NAND
  - (up to) 1000X endurance
  - **6TB per 2-socket system**
  - **Cheaper than DRAM**
  - SSDs first (demonstrated last year )
  - Intel DIMMs for next gen platform

Our focus for  
SAP HANA...



Source: [Persistent Memory: What's Done, Coming Soon, Expected Long-term](#) by Andy Rudoff, Intel

# Bringing Persistent Memory to in-memory computing

## Opportunities:

- Increased scalability
  - Larger memory modules means more memory available per server
- Significant cost savings
  - PM is cheaper than DRAM
- Improved recovery times

## Challenges:

- Higher (than DRAM) latency impacting performance
- New technology, standards still evolving...
  - Means slow, phased implementation with increased complexity and uncertain timelines

# SAP HANA focus: SNIA programming model for use with memory-like NVM

Intel DIMM based on 3D XPoint™ technology is fully aligned with the SNIA programming models for NVM and provides support for Block, File, Volume, and Persistent Memory (PM) File mode.

## Persistent Memory Mode (SNIA)

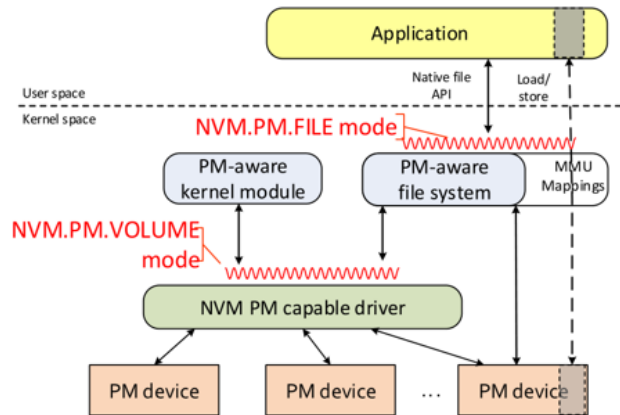
Use with memory-like NVM

### NVM.PM.VOLUME Mode

- Software abstraction to OS components for Persistent Memory (PM) hardware
- List of physical address ranges for each PM volume
- Thin provisioning management

### NVM.PM.FILE Mode

- Describes the behavior for applications accessing persistent memory Discovery and use of atomic write features
- Mapping PM files (or subsets of files) to virtual memory addresses
- Syncing portions of PM files to the persistence domain



SAP HANA focus is on leveraging memory mapping in NVM.PM.FILE mode which enables direct access to persistent memory using CPU instructions

# Sample use case: Reducing in-memory database down-time

Logging, warehousing, processing  
information: lifeline of companies



Information availability depends on  
database availability (9s)



Minimize restart time to improve  
database availability

Availability	Annual DownTime
97%	11 days
98%	7 days
99%	3 days 15 hrs
99.9%	8 hrs 48 min
99.99%	53 min
99.999%	5 min
99.9999%	32 sec

- Each restart for an IMDB can take upto 1 hour to load TBs of data to memory.
- Dell study shows millions of dollars lost per hour due to downtime\*\*
- Existing HA solutions increase the price *exponentially* for every nine

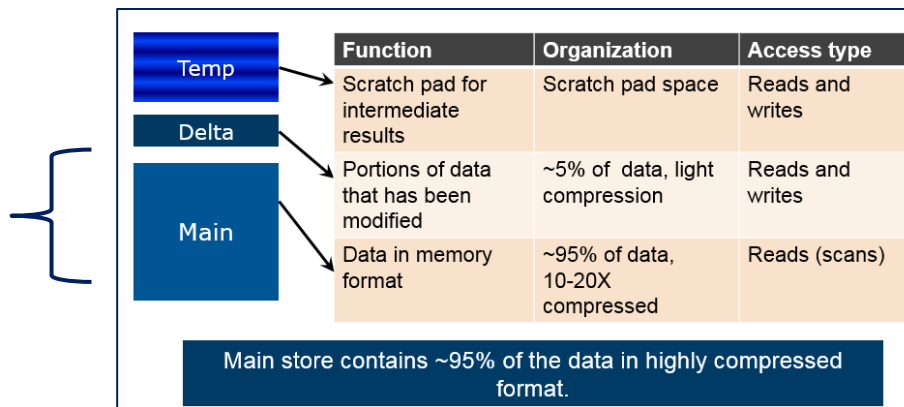
\*\*[http://tanejagroup.com/files/Compellent\\_TG\\_Opinion\\_5\\_Nines\\_Sept\\_20121.pdf](http://tanejagroup.com/files/Compellent_TG_Opinion_5_Nines_Sept_20121.pdf)



# Evolving the HANA SW Architecture to incorporate 3D XPoint™ PM technology

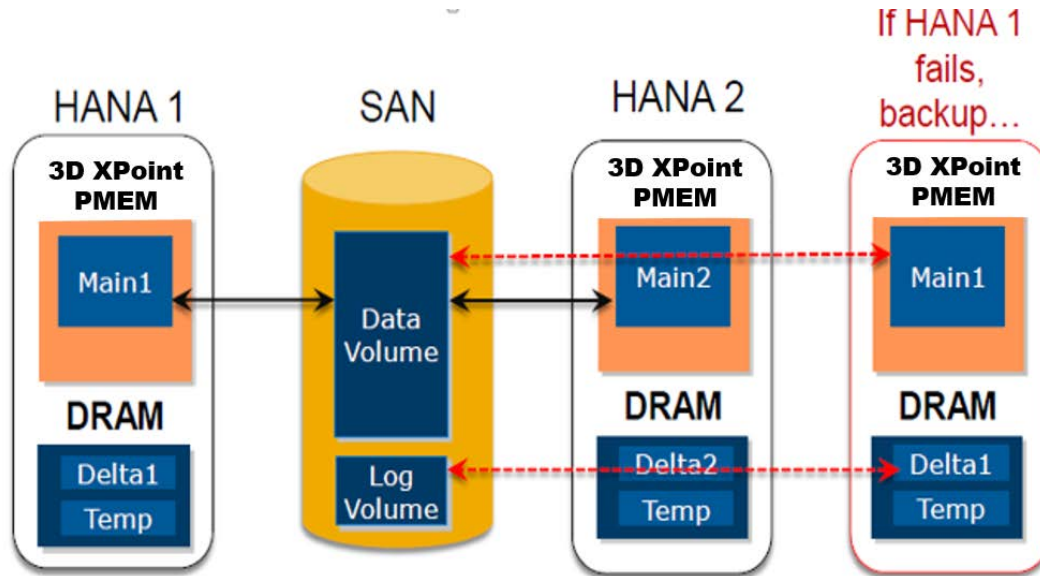
## SAP HANA: Memory Architecture

Main Store is the perfect fit for Intel DIMM based on 3D XPoint™ technology!



Technology differentiators	Why Main is well suited	
+	Large capacity	Since 95% of data contained in main, HANA can scale up to larger datasets due to increased memory capacities
+	Persistence	Avoid loading data from storage and reduce downtime
-	Higher latencies	References are in form of scans. Hardware and software prefetchers can hide latencies for such reference patterns

# Leveraging 3D XPoint™ PM technology for SAP HANA Main Store: Solution Overview



- Primary data store is data volume (in SAN or local storage)
- Main is in 3D XPoint™ PM instead of DRAM and is now persistent
- On Restart: Main already in 3D XPoint™ PM, no need to load data from SAN
- On HW failure: Backup server loads data from data volume

# Leveraging 3D XPoint™ PM technology for SAP HANA Main Store: Performance analysis

Initial results from a prototype (that leverages a suite of simulation software and hardware technologies by exploiting NUMA configuration) are very promising:

- Significant improvements in the database restart time
  - >5X improvement measured
- Acceptable performance impact
  - Persistent Memory technology has higher (than DRAM) latencies, resulting in slightly lower performance
    - Measured (simulated) performance degradation was within the expected range for most workloads

The challenge ahead:

- New technology, standards still evolving will slow implementation timelines

# Leveraging 3D XPoint™ PM technology for SAP HANA Main Store: Pros and Cons

## PROS +

- Significant improvements in database restart time
  - No need anymore to load data, Main Store is now persistent as it is in 3D XPoint™ PM instead of DRAM
- Increased memory capacity at a lower cost
  - Potentially significant cost benefits in several areas

## CONS -

- New approach has two redundant persistent copies
  - Dual-writes to 3D XPoint™ PM and “SAN”/Disks
    - Necessary for Backup and Recovery using traditional persistence; also to safeguard against corruptions in persistent memory DIMMs
- 3D XPoint™ technology is still evolving, making s/w design tricky
  - When to operate directly on “PM” vs “DRAM” ? (Data storage vs intermediate data handling)
    - Both storages co-exist in DDR4 form factor, but the capacity ratio is unknown

- A hybrid PM + DRAM approach will provide a large amounts of fast, cost-efficient memory for SAP HANA in-memory computing platform
- Persistent Memory can help with achieving near-instant recovery for main-memory databases
- Early results from bringing 3D XPoint™ PM technology to SAP HANA clearly support the expectations of persistent memory becoming a true game changer for in-memory computing

# © 2017 SAP SE or an SAP affiliate company. All rights reserved.

---

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. Please see <http://global12.sap.com/corporate-en/legal/copyright/index.epx> for additional trademark information and notices.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors.

National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP SE or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP SE or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

In particular, SAP SE or its affiliated companies have no obligation to pursue any course of business outlined in this document or any related presentation, or to develop or release any functionality mentioned therein. This document, or any related presentation, and SAP SE's or its affiliated companies' strategy and possible future developments, products, and/or platform directions and functionality are all subject to change and may be changed by SAP SE or its affiliated companies at any time for any reason without notice. The information in this document is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. All forward-looking statements are subject to various risks and uncertainties that could cause actual results to differ materially from expectations. Readers are cautioned not to place undue reliance on these forward-looking statements, which speak only as of their dates, and they should not be relied upon in making purchasing decisions.