

SAP HANA Platform

# Top Nine Questions for Choosing In-Memory Databases

SAP



# Introduction

As an experienced customer with a range of SAP solutions, you know that enterprise applications are in our corporate DNA. The SAP HANA platform not only delivers a modern in-memory computing platform that integrates with your existing IT landscape. It also provides an enterprise-grade platform to build your next-generation business applications.

The SAP HANA platform is based on breakthrough in-memory computing technology that enables rapid data access to and deep analysis of live data. Unlike other vendors who merely add in-memory cache to a disk-based database, SAP HANA is 100% in-memory computing from the ground up. Its in-memory columnar databases eliminate disk latency and

enable the processing of live transactional and analytical data on a single data copy, yielding the real-time insights your users need to act in the moment and turn your data into a competitive advantage.

Available on premise or in the cloud, SAP HANA can support business and IT transformation through agile pilots that test new data-driven tools, like automation, machine learning, and app development, in a controlled and cost-effective environment.

To help confirm SAP HANA is the right platform choice for you, here are the top 9 questions you should ask when considering your in-memory computing database.



# Contents

1. Are my applications accelerated without manual intervention and tuning?
2. Can I achieve predictable response times for ad hoc queries?
3. Can I get the full picture of my business in real time?
4. Can I answer all my data discovery questions without adding more DBAs to do data tuning?
5. Can my application provide analytics while updating the same copy of data in real time?
6. Am I able to run my business in real time, with all my data in-memory, ready for processing?
7. Can I simplify my IT landscape with an in-memory solution?
8. Can I run transactions and analytics on the same system without adding more DRAM and CPU resources?
9. Does my database vendor have an in-memory database designed specifically for both transactions and analytics?

Find Out More





# 1. Are my applications accelerated without manual intervention and tuning?

## Yes

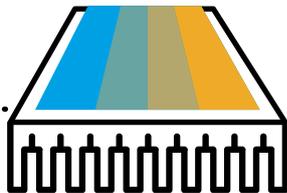
### SAP HANA Platform

Because SAP HANA stores data in-memory first.

All data is in-memory

No tuning necessary

Only 1 copy of data



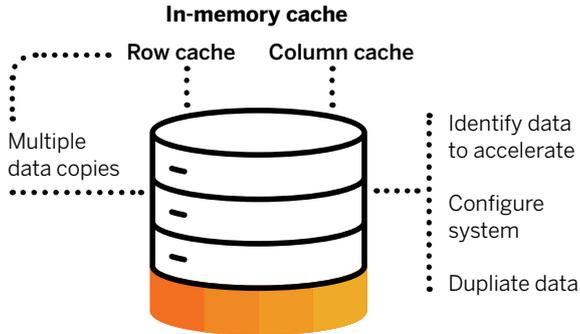
No configuration needed



## No

### Disk-Based Database with In-Memory Cache

Because all data is on the disk by default.



Manual configuration needed



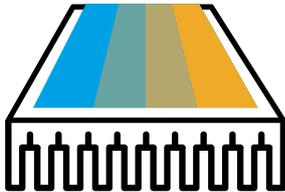
## 2. Can I achieve predictable response times for ad hoc queries?

### Yes

SAP HANA Platform

Because all hot data is automatically in-memory.

All hot data is in-memory



Predictable response times



### No

Disk-Based Database with In-Memory Cache

Because only pre-selected data is copied in-memory.

In-memory cache  
Row cache    Column cache

Constant in-memory cache synchronization



Only this data is in-memory

No predictable response times



# 3. Can I get the full picture of my business in real time?

## Yes

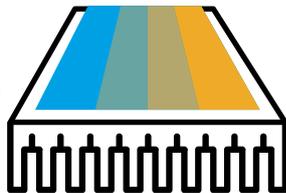
SAP HANA Platform

Because all hot data is automatically in-memory.

All data is in-memory

No limit to the granularity

Same copy of data for transactional and analytical applications



Full business picture is available in real time



## No

Disk-Based Database with In-Memory Cache

Because only pre-selected data is copied in-memory.

In-memory cache  
Row cache    Column cache



Need database administrator (DBA) to identify and physically copy all needed data to the in-memory cache

No full picture - advanced knowledge of the drill-downs is needed



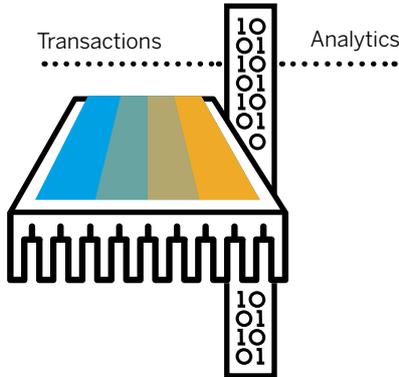


# 5. Can my application provide analytics while updating the same copy of data in real time?

## Yes

### SAP HANA Platform

SAP HANA is a true in-memory database that performs both transactions and analytics.



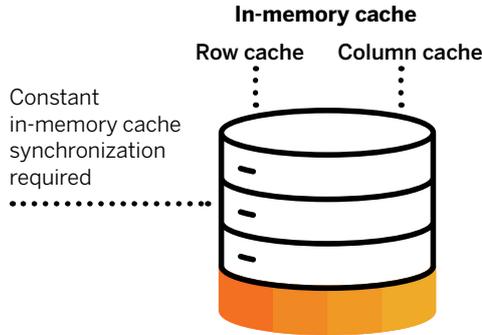
Transactions and analytics performed on same copy of data in-memory



## No

### Disk-Based Database with In-Memory Cache

The in-memory cache is not a database, it is just a read-only cache.



Transactions are performed only on data on the disk

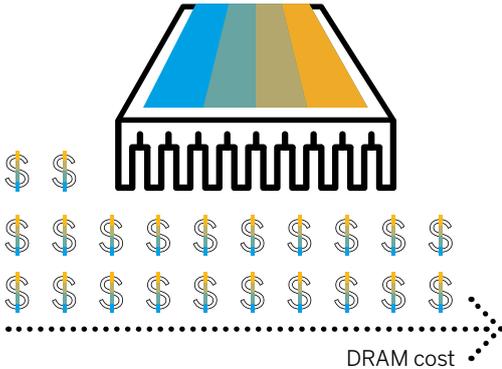


# 6. Am I able to run my business in real time, with all my data in-memory, ready for processing?

## Yes

SAP HANA Platform

Because SAP HANA is designed to manage data in-memory.



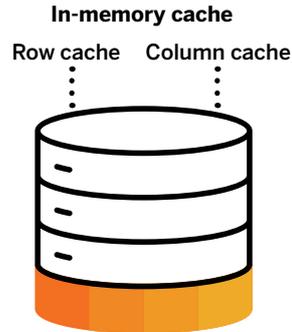
Future-proof - designed to run business at the speed of memory



## No

Disk-Based Database with In-Memory Cache

Because the solution is designed to manage data on disk.



Business innovation is limited by the speed of the disk

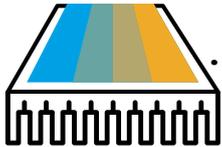


# 7. Can I simplify my IT landscape with an in-memory solution?

## Yes

### SAP HANA Platform

Because SAP HANA is an all inclusive platform.



Eliminates aggregates and indexes

Avoids data duplication for operational reporting

Provides an integrated system for any data type and data processing

Encompasses application server and advanced analytics

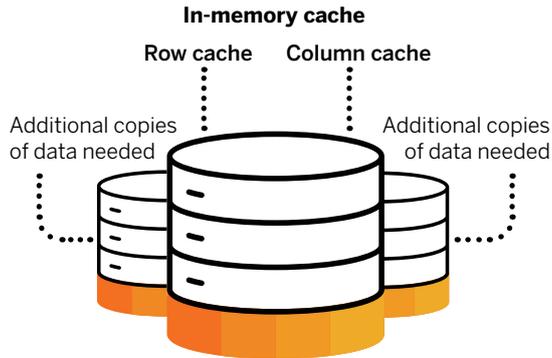
One copy of data for all requests



## No

### Disk-Based Database with In-Memory Cache

Because it is only a database.



Multiple copies of data needed for different requests

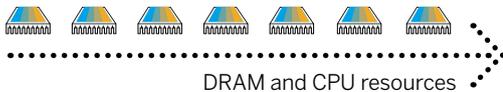
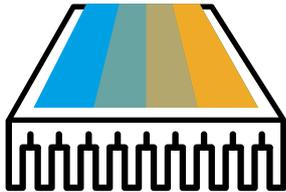


# 8. Can I run transactions and analytics on the same system without adding more DRAM and CPU resources?

## Yes

### SAP HANA Platform

Because SAP HANA requires only one copy of data.



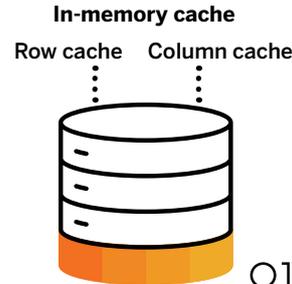
No additional DRAM/CPU needed



## No

### Disk-Based Database with In-Memory Cache

Because data is copied and stored multiple times for transactions and analytics.



More DRAM and CPU resources needed if running transactions and analytics

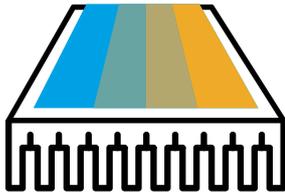


# 9. Does my database vendor have an in-memory database designed specifically for both transactions and analytics?

## Yes

SAP HANA Platform

Because SAP HANA is an in-memory database proven for both transactions and analytics.



**5,800+**  
Customers and  
**1,800+**  
Startups using it  
since 2010



## No

Disk-Based Database with In-Memory Cache

No plans announced



## Find Out More

### Learn More About SAP HANA

Check out the [website](#) which has valuable resources for fast-tracking your knowledge of SAP HANA® and a rich support section designed to help you get the highest quality answers quickly and easily from SAP experts.

[Read our blog](#)  
[Developer trial](#)

### Contact a Representative

SAP is here to help. Contact your local SAP representative:

T: +1-877-727-1127 ext 11001

### Ask SAP HANA

Connect with experts via [email](#)

### Get Involved in the Discussion

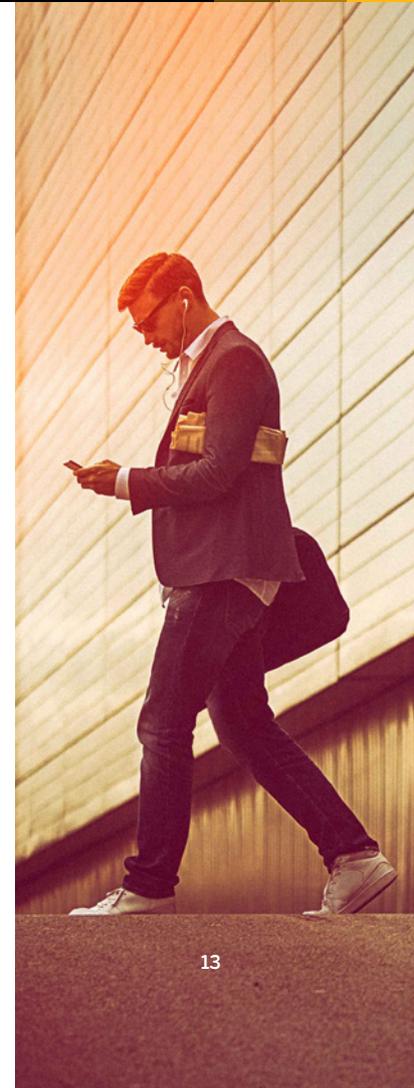
Engage with community experts on the SAP Community program to accelerate the development of HANA powered solutions.

[Engage now](#)

### Join the Conversation

Follow [@SAPinMemory](#)  
Keep updated with #SAPHANA

### Spread the Word



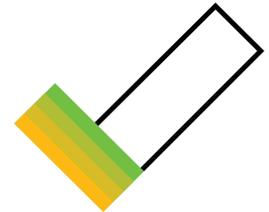
# 1. Are my applications accelerated without manual intervention and tuning?

## YES - No configuration needed

---

With SAP HANA neither specialized caches nor multiple data copies are needed. SAP HANA is architected to manage data in-memory by default and uses a columnar store to seamlessly run both analytics and transactions. Additionally, it uses CPU caches to work on compressed data, leverages multi-core processors to scan columns in parallel and uses SIMD instructions to simultaneously process multiple data sets, delivering unmatched performance without indexes or materialized views. Applications are automatically accelerated because the data they need is readily available in real time, without DBA intervention.

Close X



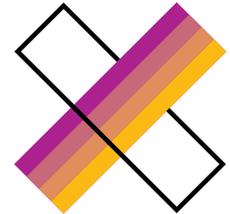
# 1. Are my applications accelerated without manual intervention and tuning?

## NO - Manual configuration needed

---

Disk-based databases are architected to manage data on disk and use add-on in-memory caches to selectively accelerate access to portions of data. To accelerate applications, DBAs identify what data to accelerate and ensure it is duplicated in the appropriate in-memory cache. Additionally they maintain indexes and other data structures to ensure that performance of both transactional and analytical workloads is acceptable. The task is daunting since it is often challenging to predict what data needs acceleration and to ensure mixed workload performance, which leads to extensive testing and tuning.

Close X



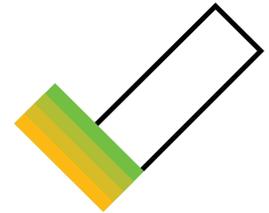
## 2. Can I achieve predictable response times for ad hoc queries?

### YES - Predictable response times

---

SAP HANA Platform maintains data in-memory by default, and all queries - planned and unplanned (ad-hoc) - have real-time access to the data they need and can return results in a predictable amount of time. It is possible to estimate response times using processors' scan speed, memory access time and size of data accessed by a query. SAP HANA also maintains one single copy of the data to be used for both analytical and transactional workloads, eliminating possible delays related to data synchronization and ensuring all queries are performed on fresh data.

Close X



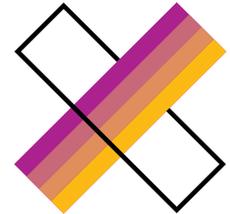
## 2. Can I achieve predictable response times for ad hoc queries?

### NO - No predictable response times

---

Disk-based databases are architected to manage data on disk and use add-on in-memory caches to accelerate access to portions of data. Only queries that exclusively access cached data can return results in a predictable time. All other queries experience delays because access times depend on parameters such as where the data is placed on disk, which indexes are used, and the number of parallel processes accessing the disk. Ad hoc queries will experience unpredictable response times when some or all of the data required is available only on disk.

Close X



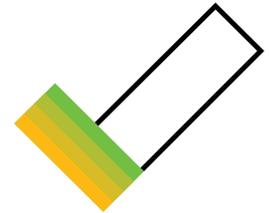
### 3. Can I get the full picture of my business in real time?

**YES - Full business picture is available in real time**

.....

Since SAP HANA Platform maintains data in-memory by default, data can be aggregated on-the-fly along any dimension, without requiring indexes, pre-aggregates or materialized views. As a result, SAP HANA can not only return aggregates in real time but also allow users to drill down to any level of detail to analyze data. With SAP HANA, users can analyze data at any level of granularity in a self-service manner and obtain results in real time.

Close X



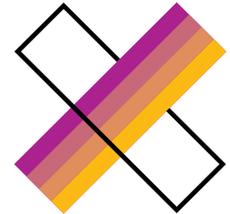
### 3. Can I get the full picture of my business in real time?

**NO - No full picture - advanced knowledge of the drill-downs is needed**

.....

Users are limited in the granularity of their real-time data analysis since disk-based databases are architected to manage data on disk and use a combination of add-on in-memory caches, indexes, pre-aggregates and materialized views to accelerate performance. For better performance the DBA has to duplicate data on in-memory caches and create data structures for data on disk before users can perform drill-downs. Additionally, if the size of the in-memory cache cannot accommodate all the required data, the DBA has to resize the system and potentially restart it.

Close X



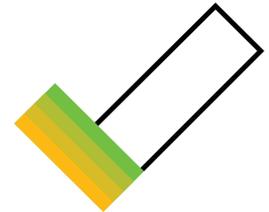
## 4. Can I answer all my data discovery questions without adding more DBAs to do data tuning?

**YES - No additional DBA time needed**

.....

SAP HANA Platform is architected to manage data in-memory by default. Applications are automatically accelerated because the data they need is readily available in-memory. Additionally, both planned and unplanned questions are answered in real time and drill-downs along any dimension are possible without additional DBA intervention. No data copies, indexes, pre aggregates or materialized view are required to deliver real-time performance.

Close X



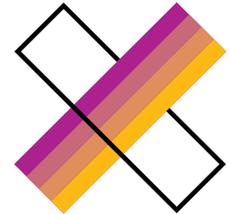
## 4. Can I answer all my data discovery questions without adding more DBAs to do data tuning?

**NO - More questions = More DBA time**

.....

Disk-based databases are architected to manage data on disk and use add-on in-memory caches and other data structures to accelerate performance. Before users can analyze data along new dimensions, DBAs need to configure and tune the database to ensure acceptable response times. This might involve copying data to the in-memory cache, dropping or creating new indexes, or creating materialized views. Some of these actions might negatively affect the performance of other applications and thus require additional tuning. For example the time required to update indexes can slow down transactional applications.

Close X

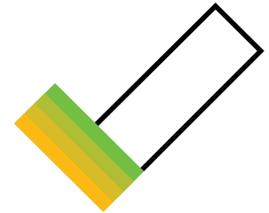


## 5. Can my application provide analytics while updating the same copy of data in real time?

**YES - Transactions and queries performed on same data in-memory**

.....

SAP HANA Platform is an ACID, persistent, in-memory, columnar database that accelerates both queries and transactions using one data copy, in-memory. SAP HANA column table's temporary delta store makes it efficient to process high-speed transactions. With SAP HANA Platform, applications can execute transactional and analytical workloads in parallel while preserving data integrity and system performance.



Close X

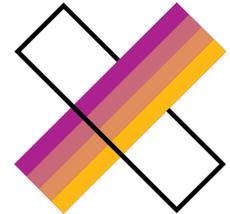
## 5. Can my application provide analytics while updating the same copy of data in real time?

NO

.....

Disk-based databases with an in-memory cache typically process transactions (update, insert or delete operations) on disk and then update the data in the in-memory cache to ensure read consistency. Applications that perform both transactions and queries would have to wait for transactions to complete on disk before accessing the updated data in the cache. Since disk access is much slower than memory access these applications will experience delays.

Close X



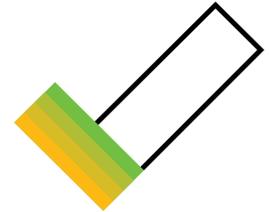
## 6. Am I able to run my business in real time, with all my data in-memory, ready for processing?

**YES - Future-proof - designed to run business at the speed of memory**

.....

SAP HANA in-memory platform maintains one data copy for both transactional and analytical workloads. All data is in a compressed, columnar format to maximize access speed and the amount of data managed in-memory. Using the dynamic tiering capability, rarely accessed data can also be maintained on disk-based, columnar tables. Access to this data remains fast because it can be moved to memory for processing without being reorganized. In this way, SAP HANA can manage databases of any size without being limited by the amount of available memory in a system.

Close X

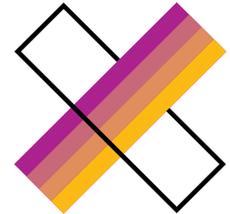


## 6. Am I able to run my business in real time, with all my data in-memory, ready for processing?

**NO - Business innovation is limited by the speed of the disk**

.....

Disk-based databases are architected to manage data on disk and use caches and other data structures to accelerate data access. In this way the disk latency can be mitigated but not eliminated because the core engine has been optimized to manage data that reside on disk and cannot function if the data is not maintained in it. While analytical workloads on cached data can be processed without accessing the disk, transactional workloads always require disk access.



**Close X**

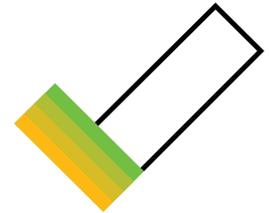
## 7. Can I simplify my IT landscape with an in-memory solution?

### YES - One copy of data for all requests

---

SAP HANA Platform simplifies IT landscapes by taking advantage of in-memory computing and by delivering application, database and integration services in one platform. By taking advantage of in-memory computing, it can efficiently process transactions, streams, graphs and advanced analytics, such as predictive, spatial and text, on the same system and on one copy of the data. Additionally, by delivering application, database and integration services in one platform, it reduces data movements and staging among operational systems and between database and application server. This results in better performance, a simplified IT infrastructure and lower administration costs.

Close X



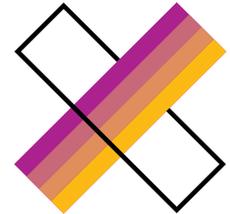
## 7. Can I simplify my IT landscape with an in-memory solution?

### NO - Multiple copies of data needed for different requests

.....

Since disk-based databases are designed to manage data on disk, in-memory caches are deployed to accelerate data access. This involves the introduction of an additional technology layer as well as the synchronization and maintenance of multiple copies of data. As a result, the consumption of system resources, the burden of system administration and the complexity of the IT infrastructure increases.

Close X



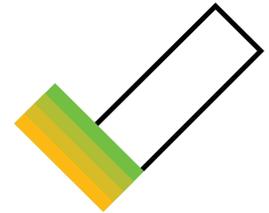
## 8. Can I run transactions and analytics on the same system without adding more DRAM and CPU resources?

**YES - Because SAP HANA requires only one copy of data**

.....

SAP HANA Platform is designed to take advantage of the latest hardware innovations. It utilizes SIMD instructions, advanced parallelization with multi-core processors and data compression to maximize CPU and RAM utilization. This allows for processing mixed workloads on the same system and on the same copy of data with the most efficient utilization of system resources.

Close X



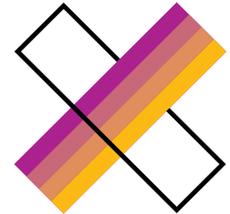
## 8. Can I run transactions and analytics on the same system without adding more DRAM and CPU resources?

**NO - More DRAM and CPU resources needed if running transactions and analytics**

.....

To process transactional and analytical workloads on the same system, disk-based databases need to use specialized in-memory caches. To leverage this new technology layer the database not only has to create multiple data copies and keep them synchronized, but in addition must also route incoming requests to the appropriate data copies. This increases the amount of CPU and memory required and delays system response.

Close X



© 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://www.sap.com/corporate-en/legal/copyright/index.epx#trademark> 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.

