

A Forrester Consulting
Thought Leadership Paper
Commissioned By Aerospike
July 2017

Hybrid Memory Architecture Drives Real-Time Systems Of Engagement

Replace Traditional In-Memory Architectures To
Better Win, Serve, And Retain Customers

Table Of Contents

- 1 Executive Summary
- 2 Hybrid Memory Architecture: A New Way Of Thinking
- 3 Memory Architecture For Databases Is A Top Priority For Businesses
- 5 Reliability And Uptime Must Be Addressed
- 7 Customers And Employees Benefit From Hybrid Memory Architecture
- 9 The Market Is Primed For The Adoption Of Hybrid Memory Architecture
- 10 Key Recommendations
- 11 Appendix

Project Director:

Sarah Brinks,
Market Impact Consultant


Contributing Research:

Forrester's Infrastructure &
Operations research group

ABOUT FORRESTER CONSULTING

Forrester Consulting provides independent and objective research-based consulting to help leaders succeed in their organizations. Ranging in scope from a short strategy session to custom projects, Forrester's Consulting services connect you directly with research analysts who apply expert insight to your specific business challenges. For more information, visit forrester.com/consulting.

© 2017, Forrester Research, Inc. All rights reserved. Unauthorized reproduction is strictly prohibited. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change. Forrester®, Technographics®, Forrester Wave, RoleView, TechRadar, and Total Economic Impact are trademarks of Forrester Research, Inc. All other trademarks are the property of their respective companies. For additional information, go to forrester.com. [1-13OB9UN]



Forrester defines hybrid memory architecture as architecture that combines the use of DRAM and SSD to intelligently move data between the two layers. It provides automation to simplify the access and processing of data, and it also offers support for transactional, operational, and analytical workloads.

68% of companies are currently using some form of in-memory architecture, and 28% are considering or planning on adopting it.

Executive Summary

In-memory platforms have become critical to build and deliver systems-of-engagement applications and support digital business transformation initiatives. However, organizations are realizing that traditional memory architectures, such as distributed caching, are failing to meet new business requirements, especially in the areas of consistency, performance, and reliability.

Hybrid memory architecture is a new approach that leverages both volatile memory, such as dynamic random access memory (DRAM), and nonvolatile memory, such as solid-state drive (SSD) and flash, to deliver consistent, trusted, reliable, and low-latency access to data in order to support existing and new generations of transactional, operational, and analytical applications. Implementing this type of architecture allows organizations to move from a two-tier in-memory architecture to a single-tier structure that simplifies the movement and storage of data without requiring a caching layer. Early adopters are seeing several benefits to their businesses, like lower costs of ownership, huge reductions in their server footprints, simplified administration, and improved scalability.

Although current hybrid memory architecture adoption is small, it is becoming an important architecture that organizations are starting to leverage. The market is seeing an increase in investment with memory architectures as new and modern systems-of-engagement applications roll out. Organizations should revisit their data management strategies for both new and existing databases to ensure that their memory architecture is delivering the reliability, uptime, consistency, and performance needed to support enterprise-wide deployment.

In February 2017, Aerospike commissioned Forrester Consulting to evaluate current levels of hybrid memory awareness and implementation within North American organizations and to determine the challenges these organizations face with their current in-memory architectures.

KEY FINDINGS

- › **Organizations face challenges with current in-memory architectures.** The day-to-day burdens of current in-memory architectures, such as high cost of ownership, issues with reliability and uptime, and an increasing number of databases, have left the market open to new architecture options that ease daily challenges and better fit business and technology needs.
- › **Many companies are planning to expand their memory-based database architecture investments to support growing business needs.** Decision makers in charge of systems of engagement are planning to deploy and expand their memory-based database architectures, as well as moderately increase spending, over the next one to two years.
- › **Hybrid memory architectures deliver better consistency, reliability, and low-latency access to active data.** Companies across a range of industries identified key use cases where hybrid memory architecture better met their needs than traditional in-memory architecture. They saw improvements to key challenges like reliability and uptime, which led to better performance and higher satisfaction than with traditional in-memory architectures.


Hybrid Memory Architecture: A New Way Of Thinking

Although in-memory has been around for decades, recent innovation around scale-out architecture, use of SSD and DRAM in a unified approach, a higher degree of automation, and support for multi-workloads have made new architectures more doable. One of the architectures that has evolved more recently is the “hybrid memory architecture.” While organizations have yet to align to a definition of hybrid memory architecture, we find that many that have implemented such architectures have a consistent viewpoint.

Hybrid memory architecture is different from traditional in-memory architecture as it leverages both volatile memory, such as DRAM, and nonvolatile memory-based storage, such as flash or SSD (see Figure 1). The use of volatile and nonvolatile memory enables low-latency processing of large data sets, provides better reliability, and reduces total cost of ownership.

Companies today are seeing greater interest in hybrid memory architecture for various reasons including:

- › **Traditional in-memory architecture lacks predictable performance.** Companies that want to be leaders in their industry and outperform their competition need to break free from legacy architectures. Traditional in-memory architecture relies on a two-tier system with a caching layer between operational databases and applications; this leads to a high cost of ownership, latency issues, and scalability issues. No one needs yesterday’s data tomorrow. Hybrid memory architecture removes the cache layer from the technology stack, thus simplifying the process, improving speed and performance, and minimizing the total cost of ownership (TCO).

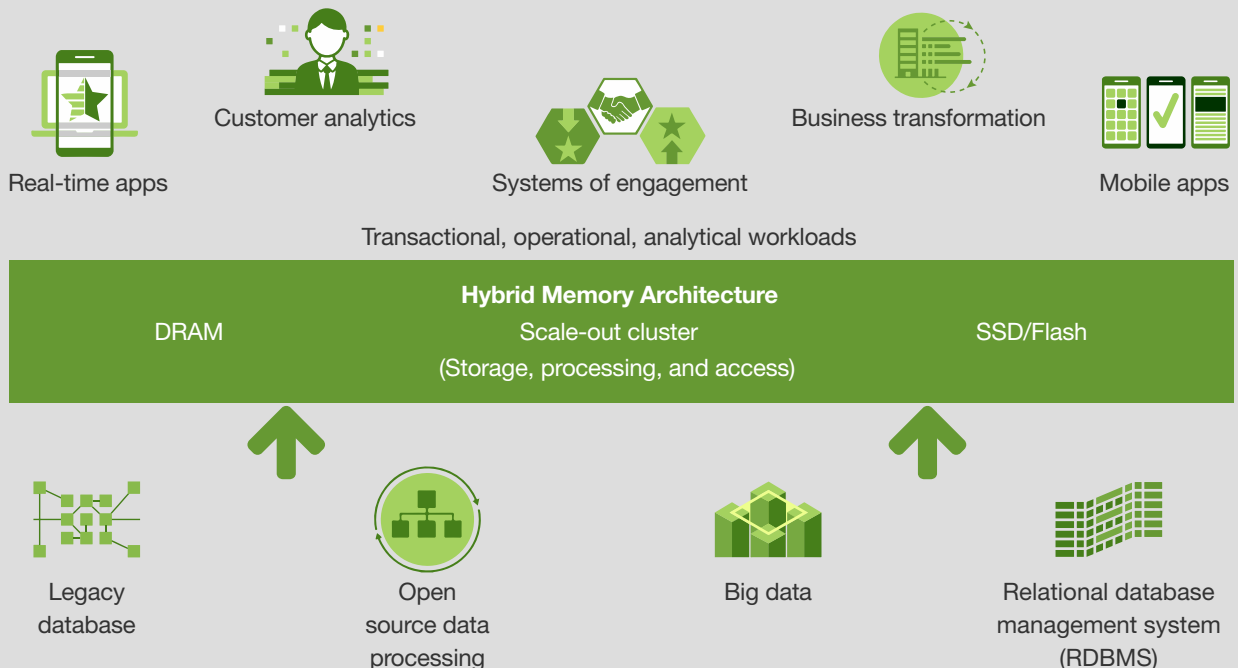


“We measure the effectiveness of the decision that we are making about in-memory platforms on three dimensions: No. 1 is speed, No. 2 is the amount of data that you can write, and No. 3 is cost. The reason cost is very important is you don’t want to spend a dollar to save \$0.05. By making it cheaper to do risk assessments, it does help us to widen the net to catch the frog, so to speak.”

Senior directors at an American online payments company



Figure 1



Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospice, February 2017

› **Organizations have found new ways to improve business.**

As businesses grow and evolve faster than ever, they need their technologies to evolve and support them equally fast. Financial services is an example of an industry where seconds can cost millions of dollars. The vice president of a brokerage and banking company recognized the need for a solution to speed up the company's trading process. Three years ago, the company was processing all its trade orders on a mainframe. This organization needed an architecture that allowed it real-time access to the stock market prices as well as its clients' data. That meant the company was processing about six terabytes of data during trading hours. It needed an architecture that could read/write at low latency, high scale, and high volume. Hybrid memory architecture was implemented, and the organization was able to completely disconnect from the mainframe during trading hours by processing everything with hybrid memory. The company could absorb the significant read/write load and greatly reduced its server footprint. It was also able to build a unique process design wherein the data was reconciled at the end of the day in batches and synced. The company has notably minimized its mainframe cycles for a competitive edge.

HYBRID MEMORY ARCHITECTURE CAN HELP REDUCE SERVER SPRAWL AND SIMPLIFY ADMINISTRATION

A key cost-saving benefit of hybrid memory is the ability to reduce the server and storage footprint for databases. Unlike traditional caching architectures that often duplicate data from transactional systems and store copies across caching servers, hybrid memory architectures don't need caching servers. This new cacheless architecture helps remove the caching layer from the technology stack; this not only lowers TCO by way of server reduction and simplifies administration with fewer number of servers, but also delivers accurate and consistent real-time information with low-latency access (see Figure 2).

In addition, hybrid memory architectures offer companies the following key capabilities:


- › **Ability to support mixed workloads.** Hybrid memory architectures deliver a modern platform to support any new business application. As the volume of data that companies use to support their real-time applications grows, the flexibility that hybrid memory offers is critical to drive business innovation and growth.

- › **Increased speed and reduced server footprints.** The vice president of a brokerage and banking company ran more than 900 million read and write transactions during trading hours using traditional caching architecture. The high cost associated with that level of volume drove the organization to adopt hybrid memory architecture, saving millions of dollars by bypassing the mainframe and processing transactions right from memory. Senior directors at an online payments company and the CTO of a mobile marketing platform were able to leverage hybrid memory to increase their server storage capacities and reduce their number of servers by a factor of four.
- › **Automation and speed.** Hybrid memory architecture helps deliver automation and simplification, reducing the effort needed to manage in-memory data platforms. The vice president of a brokerage and banking company stated: “It’s not only the server cost but also the people on the infrastructure team managing them. They take care of the data availability and the data consistency. All of that was simplified, and the operational costs were reduced a lot in terms of total cost, not just the hardware.” Interviewees for this study were able to reduce their database servers by a factor of four to 20, depending on the company and use case.

Figure 2

Use case example of server reduction:

Traditional in-memory architecture = 400 servers



Hybrid memory architecture = 20 servers



Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017
 Interview with the vice president of a brokerage and banking company, February 2017

THE NEW HYBRID MEMORY ARCHITECTURE DELIVERS INNOVATION AND GROWTH

Forrester defines hybrid memory architecture as architecture that combines the use of DRAM and SSD to intelligently move data between the two layers. It provides automation to simplify the access and processing of data, and it also offers support for transactional, operational, and analytical workloads.

Hybrid memory architecture is not just like storage or an extension; it should be treated more as a native part of that platform. It leverages SSD and flash natively in an optimized and engineered manner. Some of the key characteristics of hybrid memory architecture include:

- › Intelligently managed SSD storage; stores indices in DRAM and data in SSDs.
- › Multithreaded, massively parallel systems.
- › Reliability and uptime that is managed at the database level.
- › Predictable performance, irrespective of workload.
- › Data consistency.

Survey results show:

- › **Early adopters express high satisfaction.** Only 15% of the market has currently deployed hybrid memory architecture. Those early adopters have expressed high satisfaction. The CIO of an India-based telecommunications company described hybrid memory architectures as “a technology that really allows me to actually digitize the entire company, because it allows me to do a lot of different things in-memory, in real time. It also isolates a lot of the complexity, as well as the fragile nature of all the old systems, and frees me from all those shackles to be able to provide an exceptional customer experience. I’m able to service the customer with everything that I’m doing digitally without actually having to constantly expose myself to all these problems.”
- › **Decision makers only grasp some of the aspects of hybrid memory architecture.** The market has yet to fully align on a definition and description of hybrid memory architecture, leading to confusion, mislabeling, and a fragmented understanding of its full capabilities and benefits. Organizations understand hybrid memory architecture capabilities in pieces, such as faster processing of read and write operations, improved performance, and processing using both DRAM and SSDs or flash.



“Certain fraud patterns happen at very high velocity. The real-time analytics recognize those networks and can stop them before money starts exiting out of the system, which positively impacts customer experience and builds brand trust.”

Senior directors at an American online payments company



Memory Architecture For Databases Is A Top Priority For Businesses

Today, many organizations are spending time and money on their enterprise in-memory database initiatives. Companies are finding that more use cases are requiring in-memory technology in order to deliver better performance, scale, and low-latency access. These use cases include support for new and complex mobile, web, and interactive workloads such as real-time analytics, operational analytics, extreme transactions, real-time dashboards, fraud detection, online gaming, mobile ads, internet-of-things (IoT) apps, and customer intelligence insights.

Most organizations have memory architecture already. Based on our survey, 68% of organizations are using a form of memory architecture and another 28% are considering or planning their deployment. Forrester expects in-memory adoption to become critical for every business going forward to deliver faster analytics and insights.

Our research showed:

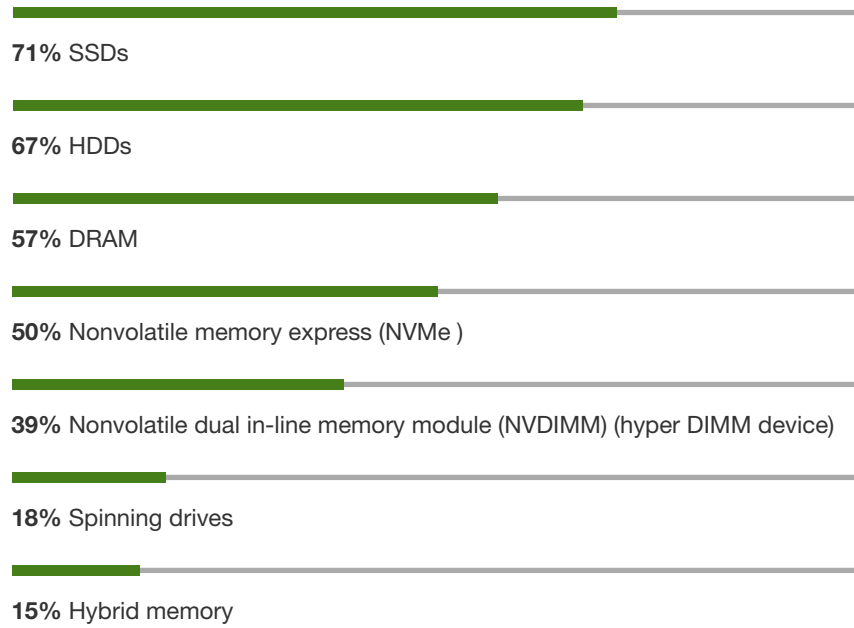
- › **Most organizations rely on traditional in-memory architecture.** Current deployment and expansion of traditional in-memory technology such as SSD and DRAM remains the highest, with 71% of organizations using SSD and 57% using DRAM. This means that most organizations today are already leveraging some sort of memory technology to support some of their business applications (see Figure 3).
- › **Larger amounts of data require real-time application support.** Based on our survey, 14% of enterprises have more than 100 TB of database size to support their real-time applications, while more than 59% have more than 60 TB of database size. With DRAM, SSD, and flash prices continuing to fall, we are likely to see petabyte-scaled deployments in the next two to three years.
- › **Architectural transformation will increase its pace.** Companies are planning to increase their spending on memory architectures as part of their digital transformation initiatives over the next one to two years. Though investment plans are modest, only 3% of respondents plan to decrease spending. Sixty-five percent of decision makers plan to increase spending by 1% to 15%, and 32% plan to spend greater than 16% over the next year or two.



57% of companies have 1,000 or more users currently using their in-memory applications.

Figure 3

“Does your company currently deploy any of the following infrastructures for memory storage for mission-critical enterprise databases?” (“Currently deployed/Expanding deployment” responses are shown)



Only 15% of companies have deployed hybrid memory architecture for their mission-critical enterprise databases.

Base: 162 North American systems-of-engagement decision makers
Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017

Reliability And Uptime Must Be Addressed

There is close alignment between the important aspects of in-memory architecture and the challenges that companies face during daily operations. Reliability and uptime are both significant challenges and are also critical aspects regarding current in-memory architecture.

Results from the study show the features that organizations consider the most important:

- › **Ninety-five percent of survey respondents listed reliability as the most important aspect of their in-memory architecture.** Sixty-five percent of respondents said reliability is very important and another 30% felt it was important, leaving only 5% who felt it was “nice to have” when it comes to their company’s digital transformation initiatives.
- › **Digital transformation initiatives are a priority, so in turn, almost all aspects of memory architecture are considered important.** Digital transformation is a cumbersome undertaking, and, as a result, nearly all performance-related aspects of a company’s memory architecture become important. Ensuring 24x7 availability and scale-out architecture is overall less important than other aspects and is rated as “nice to have.”

Results from the study show the most challenging aspects of companies’ in-memory architectures:

- › **Reliability is important but also a difficult challenge of companies’ in-memory architectures.** Companies listed reliability as their second highest challenge regarding their current memory architecture; 44% said that reliability issues are a very significant challenge for their company.
- › **Systems-of-engagement decision makers face numerous challenges.** That list of challenges mirrors what decision makers say are important aspects of their memory architecture. Reliability issues and uptime issues are some of the biggest challenges that decision makers face, and they are also some of the most important aspects of their current memory architectures.
- › **Total cost of ownership is a key area of divergence between importance and level of challenge.** Cost of ownership is the top challenge regarding companies’ current memory architectures; 49% said it was a very significant challenge. However, only 37% said that lowering cost of ownership was very important to their company, and it only ranked ninth on the list of important aspects (see Figure 4).

PERFORMANCE IS THE TOP PRIORITY AND IMPROVES WITH HYBRID MEMORY ARCHITECTURES

Organizations are supporting real-time apps where every second, and in some cases millisecond, counts. For such apps, cost is usually not an issue, since the businesses have to function at extreme speeds or potentially lose customers. Even though organizations list database total cost of ownership as their top challenge, they consider extreme performance critical for success. Our survey showed that organizations are processing enormous quantities of data with their database environments, some upwards of 100-plus terabytes, which is being used by thousands of users. Processing so much data and maintaining that much storage comes with a high cost, but delivering performance is equally important at such volume.

“Client experience is a key business benefit. We buy and sell very fast. Every few milliseconds is important to our clients. We are proactively calculating and sending, pushing information to our clients, that has improved the client experience a lot and increased our number of clients.”

Vice president, US brokerage and banking company



Our survey found that organizations are not willing to sacrifice performance to save a little. Senior directors at an online payments company experienced a shift in processing needs from 250 gigabytes to 20 terabytes of cache. Relying on a completely traditional in-memory solution to process that much data was cost-prohibitive. Incorporating hybrid memory architecture, allowing them to choose which data remained on disks and which was available in-memory, gave the directors the flexibility to maintain speed while overall reducing costs. As they said, “You don’t want to spend a dollar to save five cents.” Survey respondents listed reliability, data consistency, and improved uptime as more important than lowering the cost of ownership.

Users of traditional in-memory architectures listed a range of challenges associated with their current systems. Hybrid memory architecture addresses many of those key challenges when it is adopted.

- › **Real-time performance is the top technology benefit.** Nearly half of respondents said that improving their real-time performance would be the top benefit of hybrid memory architecture. Additional metrics that survey respondents listed as top benefits included improved uptime (41%), predictable performance (40%), and availability (37%). Uptime, predictability, and availability all roll up to improve overall memory performance.
- › **Replication is an additional performance benefit.** Though many companies are able to reduce their server footprint, they still need data to move quickly across their data centers. Hybrid memory architecture allows for very fast read and write capabilities, which also means that replication is fast. Should they experience an outage of one of their servers, the data is replicated on another server, and the company is not losing money while repairs are underway. The CTO of a mobile marketing platform uses his hybrid memory architecture to create staging environments that replicate the full environments so that when they want to run an experiment, they can roll out new software or new algorithms to a portion of their users.



Hybrid memory architecture has the highest satisfaction rating over traditional systems like SSDs, HDDs, and DRAM.

87% of companies are satisfied or very satisfied with their hybrid memory architectures.

Figure 4

“How significant are the following challenges regarding your current memory architecture?” “How important are the following aspects of a memory architecture to your company’s digital transformation strategy?”

■ Very significant challenge ■ Moderately significant challenge ■ Important ■ Very important




Base: 162 North American systems-of-engagement decision makers
Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017

Customers And Employees Benefit From Hybrid Memory Architecture

Early adopters of hybrid memory architecture have seen an impressive range of business benefits after deploying their hybrid memory architecture. Customers and employees are the top beneficiaries of a memory architecture that better fits their needs. Companies still dealing with traditional in-memory systems have the awareness to see the frustrations of their employees as well the potential to better win, serve, and retain their customers (see Figure 5).

Our study found that:

- › **Customers are the key group to benefit from hybrid memory architecture.** Nearly half of the respondents believed that hybrid memory architecture would help them improve customer experience, and nearly a third thought it would help improve customer loyalty. It would also make it possible for them to reach and serve more customers faster, which cyclically also improves customer experience.
- › **Employees also experience the benefits of hybrid memory architecture.** Employees are the frontline of defense when traditional in-memory architectures fail. Hybrid memory architecture minimizes failure pain points when applied to the right use cases by improving uptime, fast replication, and reliability. The CIO of an India-based telecommunications company said that hybrid memory architecture went “viral” at the company once one group had success with it. Another group found out about the use case and started experimenting with a second use case. Then a third group found out, and now they are replacing the entire national directory to retrieve customer data more quickly.
- › **Better performance means improved revenue opportunities.** As companies struggle with high cost of ownership, a lack of reliability, and uptime issues, they are also losing out on revenue. Having a memory architecture that better fits their needs allows them to be more efficient and do more business. The CTO of a mobile marketing platform was able to shorten the time his company takes to serve an ad from 100 milliseconds to 1 millisecond, allowing him to do much more business at a faster speed.



31% of companies think that the key business benefit of hybrid memory architecture would be improved share of voice.

Figure 5

“What would be the business benefits of a hybrid memory architecture that better fits your company’s needs?”



Base: 162 North American systems-of-engagement decision makers
Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017



One-third of companies think that an additional key business benefit of hybrid memory architecture would be improved brand perception.

The Market Is Primed For The Adoption Of Hybrid Memory Architecture


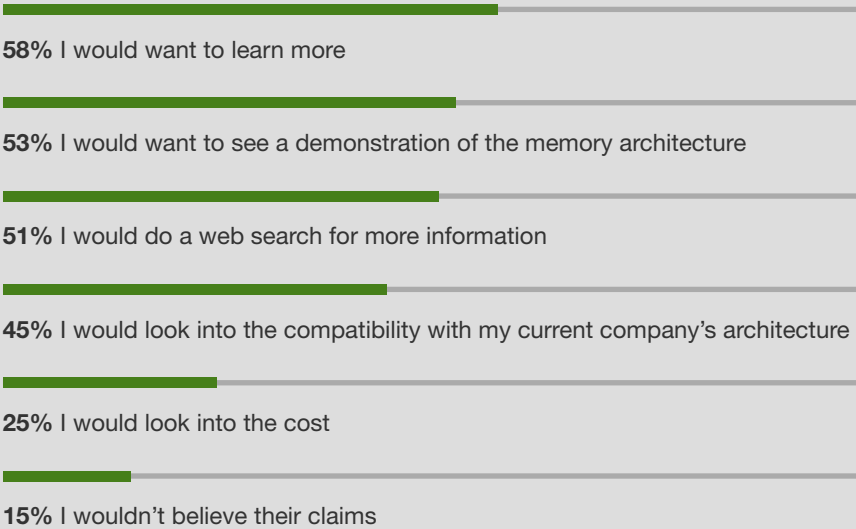
Though current adoption of hybrid memory architecture is small at only 15%, nearly half of survey respondents said they are planning to deploy hybrid memory architecture in the next 12 months. Organizations are eager to find solutions that address their list of challenges, such as improved uptime, more predictable performance, and latency.

Results from the study show:

- › **Organizations want improved in-memory architecture.** Fifty-eight percent of respondents say that they want to hear more about memory-based architecture that better fits their business needs and would ease their current pain points. Fifty-three percent of respondents say they would want to see a demonstration of the capabilities of the improved in-memory architecture.
- › **Some companies would start with their own searches.** More than half of survey respondents would start their information gathering with their own web searches while others would explore the compatibility with their current architectures (see Figure 6).

Figure 6

“If you were shown a new kind of memory architecture that ensures increased uptime, predictable performance, and real-time support and minimizes latency, how would you react?”



53% of companies are addressing their current in-memory challenges by replacing existing architectures with newer/ alternative solutions.

Base: 162 North American systems-of-engagement decision makers
Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017

Key Recommendations

Enterprises that invest heavily in databases with memory-based architecture to support new and emerging business applications, as well as use data to create personalized customer experiences, are more likely to grow faster than their competitors and win, serve, and retain more customers.



Focus on hybrid memory architecture that delivers better consistency, reliability, performance, and scale. Use hybrid memory to accelerate deployment of critical apps. With hybrid memory architecture, you can deploy apps that need extreme low-latency access to critical data. Use these solutions to scale out horizontally to leverage low-cost servers yet deliver the performance and scale needed.



Look for hybrid memory solutions that meet your requirements. Not all hybrid memory solutions are equal; therefore, look at solutions that deliver uptime, data consistency, reliability, and extreme performance and have lower TCO.



Use hybrid memory solutions to support multiple workloads. Hybrid memory architectures support various workloads including operational and analytical. Some organizations are running these workloads independently on separate clusters while others are running them together to deliver analytics at the speed of transactions.



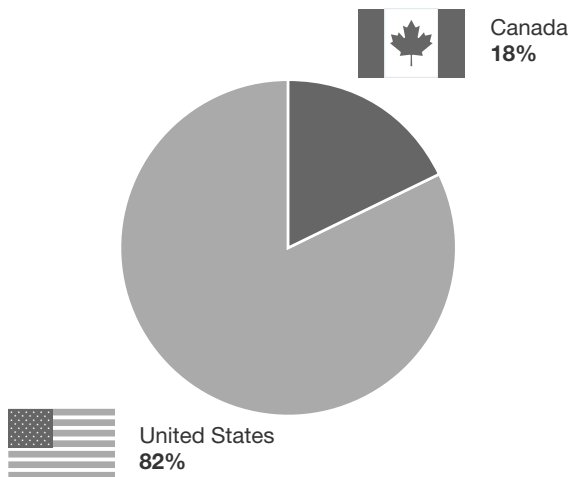
In-memory technology is mainstream; don't fall behind. Forrester estimates that 20% of all apps are suitable for in-memory databases, and this is likely to double in the next three years. Leverage in-memory to support all kinds of apps, whether they are running on-premises, in the cloud, or in a hybrid environment.

Appendix A: Methodology

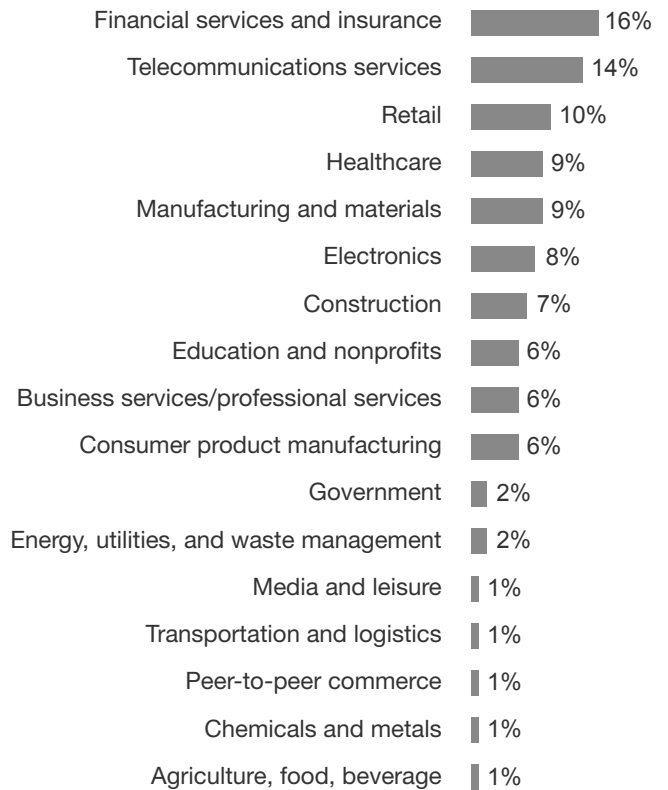
In this study, Forrester interviewed four systems-of-engagement decision makers and conducted an online survey of 162 North American systems-of-engagement decision makers to evaluate their in-memory architecture and market awareness of hybrid memory architecture. Survey participants included decision makers in IT roles. Participants were asked about their current in-memory architecture. The study was conducted in February 2017.

Appendix B: Demographics/Data


“In which country do you work?”




“Which of the following best describes the industry to which your company belongs?”



“Using your best estimate, how many employees work for your firm/organization worldwide?”



“Using your best estimate, what is your organization’s annual revenue (USD)?”



Base: 162 North American systems-of-engagement decision makers
 Source: A commissioned study conducted by Forrester Consulting on behalf of Aerospike, February 2017