ESG RESEARCH BRIEF

# Proven Performance, Reliability, and Flexibility Matter Most in the Cloud Compute War

**Date:** June 2021 **Author:** Mike Leone, Senior Analyst

**ABSTRACT:** As organizations look to improve operational efficiency, embrace infrastructure flexibility, and achieve business agility to better address real-time requirements, they continue to look to the public cloud as a strategic investment. While most organizations leverage public cloud infrastructure to reliably support several applications and workloads critical to the business, the continued promise of better performance and lower costs in the public cloud seems never ending. Foundational public cloud compute technologies in Intel-based architectures continue to deliver on the promise of efficiency, performance, and cost savings. As new compute technologies get introduced in public cloud environments with eye-catching promises of better performance and cost savings, organizations must scrutinize those promises and consider all aspects of the application lifecycle, including operational impacts and the future impact to cross-cloud flexibility.

## Shaping a Proven and Reliable Cloud

The collaboration between hardware and software vendors, thousands of ISVs, and a growing number of open-source communities have built the cloud into what it is today: proven and reliable infrastructure that redefines operational efficiency while offering limitless scale and unmatched flexibility. The adoption of public cloud infrastructure continues to be a top priority for organizations. In fact, according to ESG research, 94% of organizations leverage public cloud services today and last year's events have all but forced organizations to accelerate adoption of cloud infrastructure even more. It is a big reason why ESG research shows the area in which the largest percentage of organizations plan to increase their spending over the next year is in public cloud services.[1]



**94%** of organizations leverage public cloud services today and it is the technology area seeing the largest increase in spend over the next year.

### The Intel-based Cloud

The underpinning of the cloud was standardized on the x86 architecture. The architecture offers proven software/hardware compatibility and has continued to deliver feature improvements over decades. Organizations have relied on and continue to rely on Intel-based architectures because millions of hours have been dedicated to properly architecting, integrating,

[1] Source: ESG Research Report, *2021 Technology Spending Intentions Survey*, January 2021.

and delivering mission critical applications with Intel-based architecture features and capabilities in mind across all data center environments. They work on-premises. They work in the public cloud. They work across public clouds. This cross-environment flexibility and compatibility is critical to organizations and IT in particular, as they continue to embrace hybrid cloud operating models. Intel-based public clouds don't require IT to port an application to a new architecture, nor extensively validate an application, nor worry about maintaining multi-architected platforms. Organizations gain the flexibility they want to be able to jump from on-premises more easily to one cloud, or from one cloud to another cloud, and know that regardless of the environment or even processor generation, their applications have a high likelihood of working. And working in a way that does not disrupt the end-user experience and ensures the consistent and predictable performance that has come to be expected.

## What about ARM?

ARM is a lightweight CPU architecture that enables the processing of smaller, less complex computer instructions. By reducing the type of instructions, it can deliver fast performance and do so while consuming far less power. This is a big reason that ARM is the de facto standard processor inside common consumer devices, like smartphones, tablets, and wearables. And the promise of fast performance and low power consumption leads to interesting cost savings opportunities. For these reasons, ARM has begun to make its way into larger devices like laptops and even into the data center. And while 44% of organizations have never considered ARM-based processors or have no plans to implement ARM anytime soon, the marketing hype is there: high performance, cost savings, power efficiency. While ESG still sees very little traction of ARM in the data center today, the marketing hype is real and must be scrutinized.

**44% of organizations have never considered ARM-based processors or have no plans to implement ARM anytime soon.**

### Validating Performance Claims

While ARM can definitively deliver power-efficient performance for certain workloads in the public cloud, its ability to effectively support mission-critical applications remains to be seen. As ARM in the data center matures and benchmarking technologies are made available to third parties and customers to conduct true performance comparison testing (i.e., apples to apples), performance claims must be scrutinized and validated. ESG believes both ARM and Intel-based architectures can deliver effective levels of performance situationally for different applications, but Intel-based architectures continue to be viewed in the market as the more performant and scalable architecture for mission-critical workloads that require ultra-low latency and high throughput, cross-environment portability, and predictable performance.

**Intel-based architectures continue to be viewed in the market as the more performant and scalable architecture for mission-critical workloads that require ultra-low latency and high throughput, cross-environment portability, and predictable performance.**

### Hidden Costs of Porting to ARM

Several aspects of porting to ARM, from initial port to ongoing maintenance can present challenges. Organizations must ensure full toolchain support, including configuration managers; deployment and observability tools; and utilities, drivers, and dependent libraries. All aspects of the stack may need to be debugged, ported, and/or back ported. If resulting changes to an open-source library or utility are not unstreamed, ongoing management, maintenance, and integration with each new release becomes a cyclical burden requiring thorough testing and validation with each change.

## Measuring Success of IT Infrastructure Changes

IT is expected to deliver the right technology and the right services to the right people at the right price and the right time. This is a big reason why achieving operational excellence continues to be to the ultimate goal of IT. As IT helps facilitate migrations to cloud environments, they must embrace new infrastructure and operating models. It remains important that IT not only understand the business goals and application requirements, but also identify ways to measure success. ESG asked organizations how they measure the success of an infrastructure platform change. Improvements to performance (57%), reliability (57%), and scalability (51%) were the top three criteria that organizations used to measure success. What matters least? Cost savings, with every other quantifiable indicator coming out ahead of cost. So, while cost of course matters, everything else matters more.

### Infrastructure Change Impacting Applications

**73%** of organizations extensively revalidate their applications when moving to different hardware.

**64%** of organizations needed to change/rearchitect applications after changing the underlying infrastructure.

**49%** of organizations discovered that new hardware could not support their application requirements.

The dynamic needs of the business continue to force ongoing infrastructure upgrades and drive the need for IT modernization. Whether looking to improve performance and scalability of an existing application or to deploy net-new modern applications infused with data and next-generation technologies like AI, data centers evolve with the business. When changes occur, developers, architects, and IT work together to validate applications and ensure business and end-user requirements are met. In fact, 73% of organizations extensively revalidate their applications when moving to different hardware, whether that be with a different vendor or underlying architecture. And the more critical to the business the application, the faster it must be revalidated.

Revalidating applications drives up OpEx costs since it takes the time of developers, architects, and IT to configure, deploy, and test the environment and applications. Reducing risk of downtime is of the utmost important for all involved. If something goes wrong, not only does it take longer to migrate an application to newer/modern hardware, but it also brings into play various cost implications associated with application availability, impact to other processes, and responsiveness to other requests. Overall, ESG research found that over the past few years 64% of organizations needed to change/rearchitect applications after changing the underlying infrastructure and 49% discovered that the new hardware could not support their application requirements. Simply changing a vendor poses considerable risk. As an example, we asked organizations if they experienced challenges when making an infrastructure change. Those that also changed their vendor were more likely (by 24 percentage points) than those that did not change their vendor to find that their hardware did not support their applications requirements.

Changing the underlying architecture from an Intel-based architecture to an ARM architecture poses greater risk than that of a vendor change. The risk of incompatibilities and hidden or unexpected porting and validation costs is something IT, developers, and architects alike want to avoid at all costs. Making this type of change today definitively will require changing an application based on the way the new architecture handles instructions. And once that change is made, a new challenge presents itself: flexibility.

**The risk of incompatibilities and hidden or unexpected porting and validation costs is something IT, developers, and architects alike want to avoid at all costs.**

### Losing Flexibility and Portability

As organizations look to embrace a multi-cloud data eco-system, the introduction of cloud-based vendor proprietary computing architectures is giving organizations new computing choices. But with those options comes a new dimension of vendor lock-in. As organizations look to develop modern applications rich in data, freedom to port the application to the right cloud based on the distributed nature of data matters. Minimizing data movement, especially as it continues to grow, is increasingly important. By embracing proprietary compute platforms, flexibility and portability is inherently lost, contradicting a core principle of cloud-based computing. Until a solve for application portability based on differing compute architectures is available, the risk of vendor lock-in will remain a top consideration.

## The Bigger Truth

As organizations embrace their data-driven, cloud-centric futures, accelerating digital transformation initiatives is all but a given. With infrastructure modernization a key pillar of digital transformation, organizations want peace of mind knowing that the ultimate infrastructure destinations for their mission critical applications are defined both today and in the future. For today, research shows that performance, reliability, and scalability matter most for the business. And while predicting future outcomes is difficult for any business, lessons from the last year show that ensuring as much consistency and predictability in application behavior critically enables the personas that matter most in IT—operations teams and developers—to do their jobs as effectively and efficiently as possible. Further, enabling organizations to achieve a new level of agility and flexibility will prove beneficial as future success will be dictated by the responsiveness to the dynamic nature of the business.

What does this mean for underlying compute architectures? While the evolution of the compute stack will impact all organizations, operational efficiency and reliability is king. And research proves that today, while attractive alternatives are emerging and being evaluated based on robust marketing engines, they are still more hype than reality. Organizations are not yet willing to make leaps of faith when it comes to their data centers and mission-critical applications, especially when they may pose greater risk than reward. So, is it worth it to change your underlying Intel-based architecture to ARM to save some money? The short answer today is no…for now.