



CASE STUDY



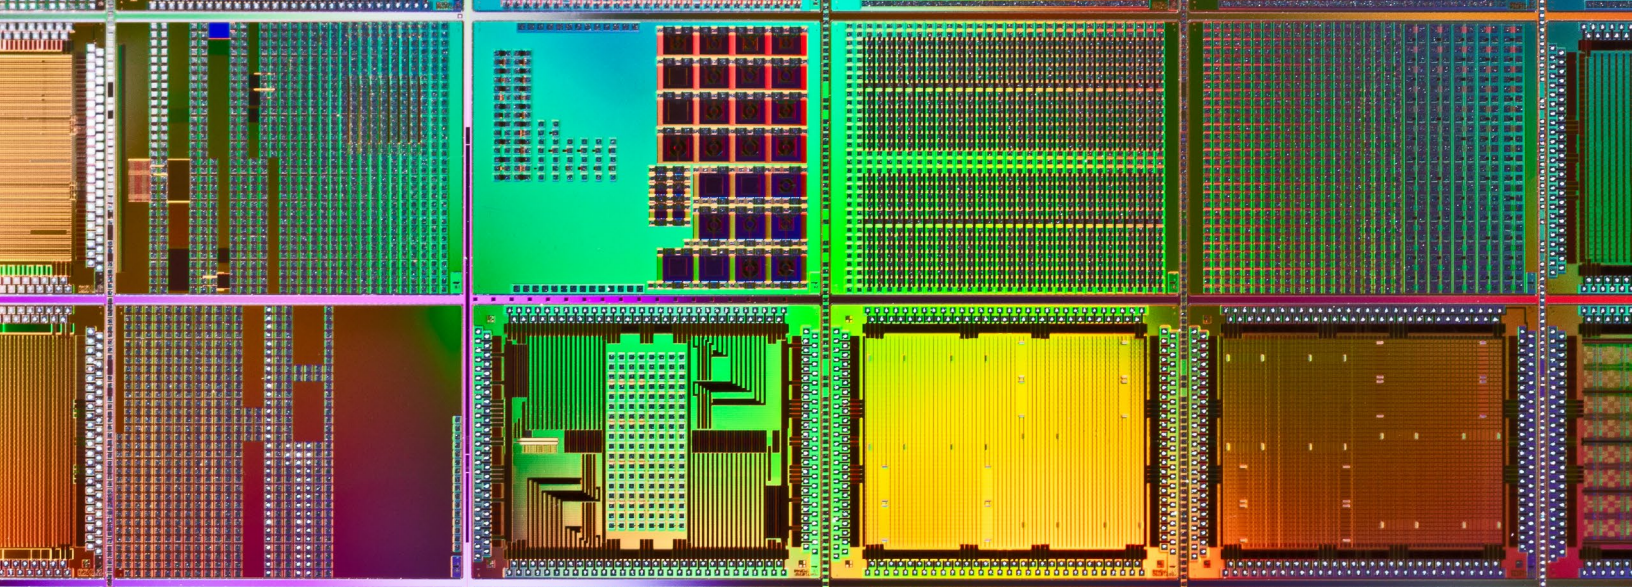
Transitioning to Arm-based Cloud: Honeycomb's Journey



TAKEAWAYS

- By transitioning to the Arm-based cloud architecture, Honeycomb has realized 50 percent lower costs and uses fewer instances without sacrificing volume or performance.
- The transition from x86 was smooth and high performance was maintained.
- Honeycomb's journey from x86 to Arm64 is a compelling example of how embracing architectures that deliver processing efficiencies can lower costs and power consumption without sacrificing performance.





Introduction

In today's rapidly evolving technological landscape, observability has become critical to understand and manage complex systems. If you're not familiar with the concept, observability refers to the ability to gain insights into the internal state of a system through by analyzing its data, including logs, metrics, and traces. In this way, engineers and operators gain a deeper understanding of system behavior, identify potential issues, and make data-driven decisions to enhance reliability and performance.

But like any critical technology, there are challenges organizations face in the observability space. One is the sheer volume, velocity, and diversity of data and alerts generated by systems. Managing and analyzing this vast amount of data in real-time demands scalable solutions that can handle the resource-intensive task effectively.

Moreover, as technology adopts distributed architectures, microservices and containers, gaining a comprehensive view of the entire application becomes increasingly difficult. The intricate interactions and dependencies between different components of the system create a complex dynamic environment that requires specialized monitoring approaches.

Fortunately, a number of companies have eagerly stepped in to provide observability services, but they too face enormous challenges as they confront the increasing complexity and volume of data that their growing lists of customers need managed. Managing and analyzing large amounts of data in real-time is incredibly resource-intensive and requires scalable solutions. Observability vendors' pricing models are based on how many logs and metrics their applications will generate from one day to the next. Yet, in some cases, observability solutions have become [so expensive](#) that most startups and medium businesses [can't afford them](#).

One leading observability platform is honeycomb.io, which helps engineering teams gain a deeper understanding of their own production systems and their end-users' experiences. Honeycomb.io has adopted a solution intended to tackle this cost burden: Shift its servers from Intel-based hardware to Amazon AWS instances based on Arm-powered AWS Graviton processors. The move has led Honeycomb to realize 50 percent lower costs and use fewer instances without sacrificing volume or performance.

Marvels of Multi-Architecture

Honeycomb's transition to Arm64 was not a rash decision but a considered experiment. Honeycomb Field CTO Liz Fong-Jones was curious as to whether Arm should be adopted within the Honeycomb infrastructure. She and her team wanted to test the waters, to compare results between Arm- and Intel-based systems safely before making any significant commitments.

They began experimenting with AWS Graviton in March 2020. Within six months, they had their first workloads in production. A year later, they had nearly all (92%) of their workloads and environments on Arm CPUs. By April 2022, they were able to turn off the last of their x86 EC2 instances and were running all of their AWS Lambda workload (except 1% for x86 compatibility validation) on the Arm instruction set.



The transition process from x86 to Arm64 was a fascinating journey for Honeycomb. As a Golang shop, they had a flexible environment that allowed cross-compiling through setting environment variables. This meant that the team could make adjustments within their build toolchain, rather than having to modify their source code. However, since they weren't containerized at that time, they had to work around operating system-based images. Fong-Jones acknowledged that the experience three years ago was more challenging because the ecosystem was less stable compared today, when such a transition would be seamless.

The smoothness of Honeycomb's transition to AWS instances can be attributed to their multi-architecture approach. By rolling back and forth between architectures, they were able to identify areas requiring additional work, such as compliance software and library dependencies. Over the past few years, Honeycomb dedicated significant effort to improve their software support and libraries. They collaborated with the community, contributed code, and benefited from the contributions of others. This collaborative effort created a more robust ecosystem and facilitated a smoother transition for everyone involved.

Honeycomb's transition to Arm64 brought several immediate and significant benefits. One of the biggest was the transition to Arm allowed Honeycomb to establish a new cost basis, which resulted in infrastructure costs that are more than 50% lower compared to those shared by large enterprises.

Honeycomb is quick to note that they don't view these beneficial dynamics as a race to the bottom of costs but rather an exercise in delivering value quickly and efficiently. They believe that as infrastructure costs fall significantly, the human resource cost becomes a more dominant factor. This means, for example, more engineers can be hired to optimize and evolve the company's products and technology.

Lower Cost, More Opportunity

Ian Smith, engineering manager at Honeycomb, [has written extensively](#) about the transition to Arm. Smith notes that Graviton has enabled Honeycomb to scale up its product without increased operational toil, spend less on compute, and have a smaller environmental footprint. Honeycomb's total cost of ownership (TCO) is among the lowest in the observability space.

“When our sales team works with large enterprises on build vs buy analysis, we find that our estimated infrastructure cost can come in more than 50% less than their ELK-based (Elasticsearch, Logstash, and Kibana) infrastructure cost they share with us,” Smith wrote. “This doesn't even include the cost of replicating our real-time querying and cutting-edge visualization technology. It is thanks to AWS innovating in what they can offer us that we can focus on our core business of providing cutting edge observability tooling to our customers.”

Exact power savings aren't known, because Amazon doesn't disclose that type of data, but the lower power consumption has allowed Honeycomb to run fewer instances. This reduction in power consumption from Arm-based Graviton chips allows Amazon to provide cost savings to customers, such as Honeycomb.io, who then have the opportunity to pass that along to their customers. Honeycomb.io also estimates lower ongoing power consumption thanks to anticipated improvements in Arm processor efficiency in future generations.

Honeycomb leverages the flexibility and optimization opportunities provided by the Arm ecosystem to deliver better experiences for their customers. For example, they recommend in some cases that customers use those AWS Graviton instances to fine-tune unique workloads — for example dynamic sampling — which can be performed cost-effectively on Arm-based instances.

Honeycomb.io has discovered another benefit: New developers building in a multi-architecture world receive Arm-powered laptops, so developing for Arm on Arm with Arm becomes a very efficient flow.

Summary

Honeycomb's journey from x86 to Arm64 serves as a compelling case study in embracing architectures that deliver processing efficiencies that can lower costs and power consumption without sacrificing performance. The company's multi-architecture approach, collaborative efforts, and contributions to the ecosystem enabled a successful transition. While challenges existed, Honeycomb's experience highlights the importance of being adaptable and supportive within a rapidly evolving tech landscape.

"We've enjoyed our journey over the past two years with the Arm and Graviton families, and we hope sharing the lessons we've learned has inspired our customers and the AWS community at large," Smith wrote. "Our fun with Arm and Graviton is far from over."

[Learn more about how to transition to Arm architecture.](#)