



Q&A

Why architecture matters

Nothing helps you discover the essence of your product more than trying to explain it to somebody else. Sheryl Koenigsberg, Director of Product Marketing at Infinio, sat down with Infinio's CEO, Arun Agarwal, to dig deeper into why customers who are considering server-side caching solutions should look closely at the solution's architecture before purchasing. A product's architecture will dictate future capabilities and associated costs - key components that can determine success or failure in a data center environment.



Q: When you say architecture matters, what do you mean?

A: Every good company has a fundamental architectural insight that drives the strategy of their product. Customers might think of products as collections of features but we need to think about what kind of framework we need to deliver that value. And choosing an architecture is what dictates the possible universe of features we can deliver and at what cost.

Q: What's a good example of that I might be familiar with?

A: Several of our colleagues were part of the team that helped build Endeca. For that product, we wanted to create a better experience for searching and navigating data. We had to build a fundamentally different database engine to provide this kind of value – and once we did, it gave us a huge competitive advantage versus the classic search engine vendors.

As an example, you're probably familiar with the Endeca technology on lots of e-commerce websites: it's the "faceted navigation", the set of filters you often see on the left side of a set of search results. Delivering this faceted navigation feature was closely aligned with how we designed the product, and it was really flexible. It may seem like our competitors could have just added that feature, but it turned out that their efforts to do so came with a huge set of caveats – like difficulty scaling or keeping up with rapidly changing data.

You just can't bolt a feature on to the wrong architecture and expect it to work.

Q: So architectural choices govern how a product develops over time.

A: Well, yes, but it's bigger than just a single product. Architectural insights that form the core of a product allow for diversification within an industry. For example, Endeca was able to take their architecture and become very successful in business analytics. Our competitors at the time went in different directions depending on their unique architectures – many of them also becoming very successful.

When industries emerge that aren't built on unique architectures, they commoditize very quickly. Obviously, at Infinio, we're interested in being in the first category.

Q: So what do you see as the “fundamental architectural insight” at Infinio?

A: One of the things we knew wanted to build was a product that customers could install and begin to use without disruption to their existing environment. We knew that leveraging RAM for this was the best choice, so we designed our cache to optimize for RAM, in small amounts, distributed across servers.

One problem is how to best use small amounts of RAM efficiently. The first thing we decided was to make a shared cache, so all the hosts in the cluster had access to all the cache. Then we looked at how to get the best deduplication rates – and that was by making the cache content-based (rather than location based). That enabled us to get inline deduplication across all the hosts and handle the vast amounts of duplicate data that is resident in virtual environments.

Those two decisions (being global and being content-based) enabled us to create the largest possible addressable cache space for all the hosts in the cluster.

For our customers, those decisions have led to SSD-like cache hit rates without the SSDs.

Q: How easy or hard is it to change architectures once the product is in the market?

A: I think it's incredibly hard. If architecture is the thing that dictates what universe of features we can deliver, changing architecture would probably mean abandoning or needing to rewrite a lot of the existing product. Once we start to have a legacy code base, it's nearly impossible to change architecture but still continue to deliver product quickly to customers.

Q: So how do we know we're making the right choice?

A: We don't. We have to make educated guesses based on what we know about the market today, and the opportunities for value we see that we can bring to our customers.

At Infinio, we optimize for the customer's experience using the product. Our VP of Product Management ran a datacenter, so he has firsthand perspective on what it's like to struggle with the things our customers struggle with. We try not to make any architectural decisions that would result in making life harder for customers.

Q: What about a more technical example?

A: Sure, on a more detailed level, we optimize for performance across our cluster. That is, we are trying to bring better performance to all Infinio-accelerated datastores equally. We use content-addressable scheme with a consistent hash to equally distribute the cache across the cluster.

Q: So if hosts go into maintenance mode or vMotion, or leave or enter the cache, all their workloads are impacted equally?

A: You got it. In fact, with vMotion, the workloads aren't impacted, their cache mapping travels with them. But for the other examples, yes, because the cache is distributed across the cluster mathematically, the impact is minimized, and lessens as the cluster gets bigger.

That's a pretty good example of an architectural decision that has palpable benefits to a customer.

This article was written by written by Sheryl Koenigsberg, Director of Product Marketing, Infinio. You can follow Sheryl on Twitter @storagediva

For more information about Infinio Systems, go to www.infinio.com, contact us at 617-374-6500, or follow us on Twitter @infinio.