

Home > AWS

Back to Las Vegas: our journey through the AWS re:Invent 2021!

12 December 2021 - 10 min. read

AWS re:Invent

A week has passed since the end of the re:Invent 2021, the most awaited moment of the year for those who, like us, work in the IT world.

This year signed the "back to physical" for the AWS re:Invent as a great technological event after exactly two years (my last event was, in fact, the re:Invent 2019).

I will not detail how the experience of attending an in-person event was pleasant from a personal point of view and how it allowed us to make the most out of every single idea offered by the conference if compared to the "virtual-only" mode. I would say the obvious.

Instead, I'd like to congratulate AWS for being the first in the "big ones" of IT to believe in such a huge event and for the great work done in organizing it in an in-person mode even with thousands of uncertainties and many organizational and logistical challenges that could have changed the plans.

Of course, compromises were necessary, such as a reduced number of people (20,000 attendees VS more than 60,000 in 2019) or limits to the organization of some specific events. Anyway, in general, it was a really successful conference everything considered. Let me say that with less than a third of the participants compared to the last time, it was far easier for me to enjoy the whole event.

As I have already said many times, the most exciting aspects of the re:Invent are not (or not only!) the contents presented in the (four!) Keynotes and during the hundreds of sessions: most of them, in fact, are available online in a few hours in articles summarizing the most important announcements, recordings of all the sessions conveniently available on YouTube (and, unfortunately, just because you think "I can watch them whenever I want" you'll never take the time to watch them actually, but that's another story ...)

For this reason, I'm not going to write yet another list of all the news announced, but I will try to tell more about the impressions and the takeaways that the conference left me as food for thoughts for 2022.

The peculiar aspect of the re: Invent is the possibility of "touching" all the things we usually interact with through a web console or a CLI. Also, you can interact and speak directly to the various products and services managers or with the speakers of the various technical sessions.



When I say "touching," I mean touching literally.

For example, what you see in this photo is one of the custom switches in use at AWS data centers. Despite the huge number of questions I was asking, an AWS staff member kindly told me everything about it. The "custom" aspect of this switch fascinated me: even if the data plane (the leftmost part with the largest metal heatsink) reaches very high numbers (more than 400 Terabit / s of aggregate bandwidth!) while using standard market components, the control plane (the card in the center, with the smallest heatsink) which is the one that actually has to manage all the calls to the VPC service and consequent configurations (for each VPC, for each AWS account, for several million customers!) is a proprietary ASIC under the Annapurna Labs brand, the company specializing in microelectronics that Amazon acquired in 2015 to move the hub of innovation at the hardware level (hence the AWS Nitro System was also born).

No "commercial" network equipment has a control plane with such high performance, simply because there was never any need on-premise! The need arose with the evolution of the Cloud, and therefore AWS was the first to "do it at home". It is unbelievable that this switch alone consumes more than 12kW and is powered directly by a proprietary backplane mounted directly in the rack cabinet. It was impossible to place such a large transformer inside it. This conversation alone is worth an intercontinental trip!

The atmosphere of innovation that you breathe (especially for those coming from outside the United States) is incredible. AWS re:Invent brings the most "smart" people of a large part of the "coolest" technology companies in the world together in a single place (from HashiCorp to MongoDB, from ARM to CloudFlare, from Slack to Netflix, including the "Italian" Kong and Sysdig, to name a few ...). And the best part of this is that you easily end up with a beer with them between one event and another of the conference! Experiences like this bring immense value to those who are lucky enough to experience them. That's why beSharp takes part at the re:Invent with a team that grows in number every year.



In my opinion, from the news point of view, this edition on the AWS re:Invent can be considered a bit weak. Probably this is due to the period of time we are living... or maybe as the Cloud market (and AWS in particular) can now be considered in its "mature" phase, we can no longer expect the same pace of innovation as a few years ago. The fact is that there were no announcements that made me "jump in the chair".

Certainly, the most exciting announcement is the serverless and on-demand mode available for four of the main services in the analytics scenario: Redshift, MSK, EMR e Kinesis data stream.



In particular, Redshift is one of the most popular AWS managed services. It allows you to reduce costs compared to commercial data warehouses (like Oracle) significantly. Like other similar services, it suffered a little from being very VM-centric... I am very curious to see in detail the on-demand scaling model and, obviously, the pricing model: I assume that for this service, too we can benefit from the architectural advances introduced by AWS with Aurora and then Aurora Serverless related to the separation between the transactional layer and the storage layer. Let's hope that, more than a year after the announcement, Aurora Serverless v2 will soon go into general availability. From the point of view of the time the newly annonced services spend in preview AWS has a lot of work to do!

An "obvious" announcement was Graviton 3 (the CPU with ARM architecture produced directly by AWS) and the related EC2 C7g instances.



Let me explain: the numbers seem particularly interesting, both from the pure performance and efficiency point of view, thanks also to the use of DDR5 memory and hardware support for calculations in half-precision for Machine Learning workloads).

I say it is "obvious" simply because as the ARM CPUs is one of the most rapidly evolving technological fronts at the moment, it provides at each iteration (usually annual) a double-digit percentage increase in performance and efficiency.

With this in mind, this announcement comes with the new EC2 Mac instances based on Apple's M1 CPU (also widely expected). After conquering the entire mobile sector in a short time thanks to energy efficiency, in the last 2 years, the ARM architecture has entered the world of personal computers and data centers (especially those of Cloud providers). This is also due to an evolutionary stagnation of the x86 architecture, particularly on the Intel front.

On the other hand, AMD has been innovating with ZEN architecture for some years. In this regard, I was happy to hear the announcement - finally! - of the "Milan" EPYC CPUs (ZEN3) on EC2 (M6a instances). In fact, if, on the one hand, ARM is conquering more and more slices of the server market, many workloads, for various reasons, are and will still be for a long time bound to x86 architectures (hopefully x86_64). In this case, EC2 instances with AMD CPU represent, in my opinion, an excellent choice, with a cost reduction for the same performance compared to the Intel counterpart, and in some cases, even some more features.

Instead, I believe that higher-level EC2-based managed services (e.g. Elasticache or RDS) in the future will increasingly be CPU Graviton-based, as in that case, the underlying CPU is transparent to the user, except in exceptional circumstances, such as those that require the use of RDS Custom.

Still on the infrastructure and hardware front:

(I think it is now clear that the "under-the-hood" part is the one I am most passionate about ;-)) worthy of note announcements are the general availability of the Outpost in "reduced" format to 1 and 2 rack - also launched exactly one year ago - and EBS, which gains two very interesting new features: the recycled bin and the Snapshots Archive, a new low cost storage tier - long-awaited! - for snapshots, probably based on Glacier.

 And finally three goodies from Peter DeSantis' Keynote: The screenshot (or presumed) of the August 2006 email with which Jeff Bezos asked DeSantis, Andy Jassy and Adam Selipsky for feedback on the beta testing of the newborn EC2 service. Shown in celebration of EC2's 15th anniversary.

Just a fun fact: Peter is the only one in this email who has not (yet?) Been CEO of AWS... Is that a good omen for him? :-D



 The preview of EC2 Trn1 instances (still in preview) based on the Tranium processor (AWS proprietary silicon, optimized for distributed training of Machine Learning models). This type of instance is designed for workloads that require a very high bandwidth, but I never expected to see what is anticipated for the specific Trn1n family: 1600 Gbps. Crazy!



A very in-depth explanation of the challenges of managing solid-state storage (SSD) at a "Cloud provider" scale to ensure consistent performance with ever-increasing customer demand. The "secret" is in the so-called Flash Translation Layer (FTL), actually the disk firmware. As already happened for the Nitro System and the switches mentioned above, here too, AWS decided at some point to start producing its own FTLs optimized for the scale, latency, performance, and functionality required by the Cloud. That's why Nitro SSD was born. Silently introduced last year, it is already at the second generation, which is the basis of the new I4 family of EC2 instances optimized for high-performance storage.



And come to think of it, there was an announcement that made me jump in my chair: the availability of OpenZFS as a file system for Amazon FSx! But I realize that very few nerds in the world would get excited about it.

I'm curious to test it as soon as possible and to understand the goodness of the implementation as a managed service by AWS of what remains the best file system around.

Also very interesting is the announcement of the Customer Carbon Footprint Tool (coming soon), a tool for customers to measure the ecological footprint of their workloads; in full consistency with the announcement of the sixth pillar of the AWS Well-Architected Framework: sustainability.



In closing, last noteworthy mentions:

- Secrets Detector per CodeGuru Reviewer, to avoid forgetting credentials in the code
- On the advanced networking front, AWS Cloud WAN, AWS Direct Connect SiteLink e AWS Private 5G
- Control Tower + Terraform, I see positively opening up to market standards
- AWS Migration Hub, to transform your application into Cloud-native, one API at a time
- Werner showing impressive numbers about IAM during his Keynote
- AWS Mainframe Modernization, so we don't miss anything

The just ended re:Invent was my 10th consecutive re:Invent, the ninth in-person in Las Vegas, the first as AWS Hero, and the fourth as AWS Ambassador. I must say that meeting with the members of these two fantastic communities was an enriching experience and a privilege (not just for the front row seats reserved during the Keynote!) I want to thank Ross, Farrah, Elaine, and Matthijs for this.



A special thank also goes to my beSharp colleagues with whom I shared this journey. In this photo, the whole team at the AWS booth, where we presented Leapp, our Open Source project!

Haven't you ever heard of it? Check out the project's GitHub repo, try it out, and let us know what you think!



I greet you with this beautiful view of the Strip - I really missed Las Vegas! - a little hint just to grow some suspense: one of the main themes I was interested in this year at the re: Invent was Vi saluto con questa splendida vista della Strip - quanto mi è mancata Las Vegas! - non senza anticiparvi, per creare un po' di suspence, che uno dei temi principali di cui mi sono interessato quest'anno al re:Invent è stato il Quantum Computing... but yes, we will talk about this on these pages, soon! Stay Tuned!

#Proud2beCloud



Simone Merlini

CEO and co-founder of beSharp, Cloud Ninja and early adopter of any type of * aaS solution. I divide myself between the PC keyboard and the one with black and white keys; I specialize in deploying gargantuan dinners and testing vintage bottles.

Copyright © 2011-2021 by beSharp srl - P.IVA IT02415160189