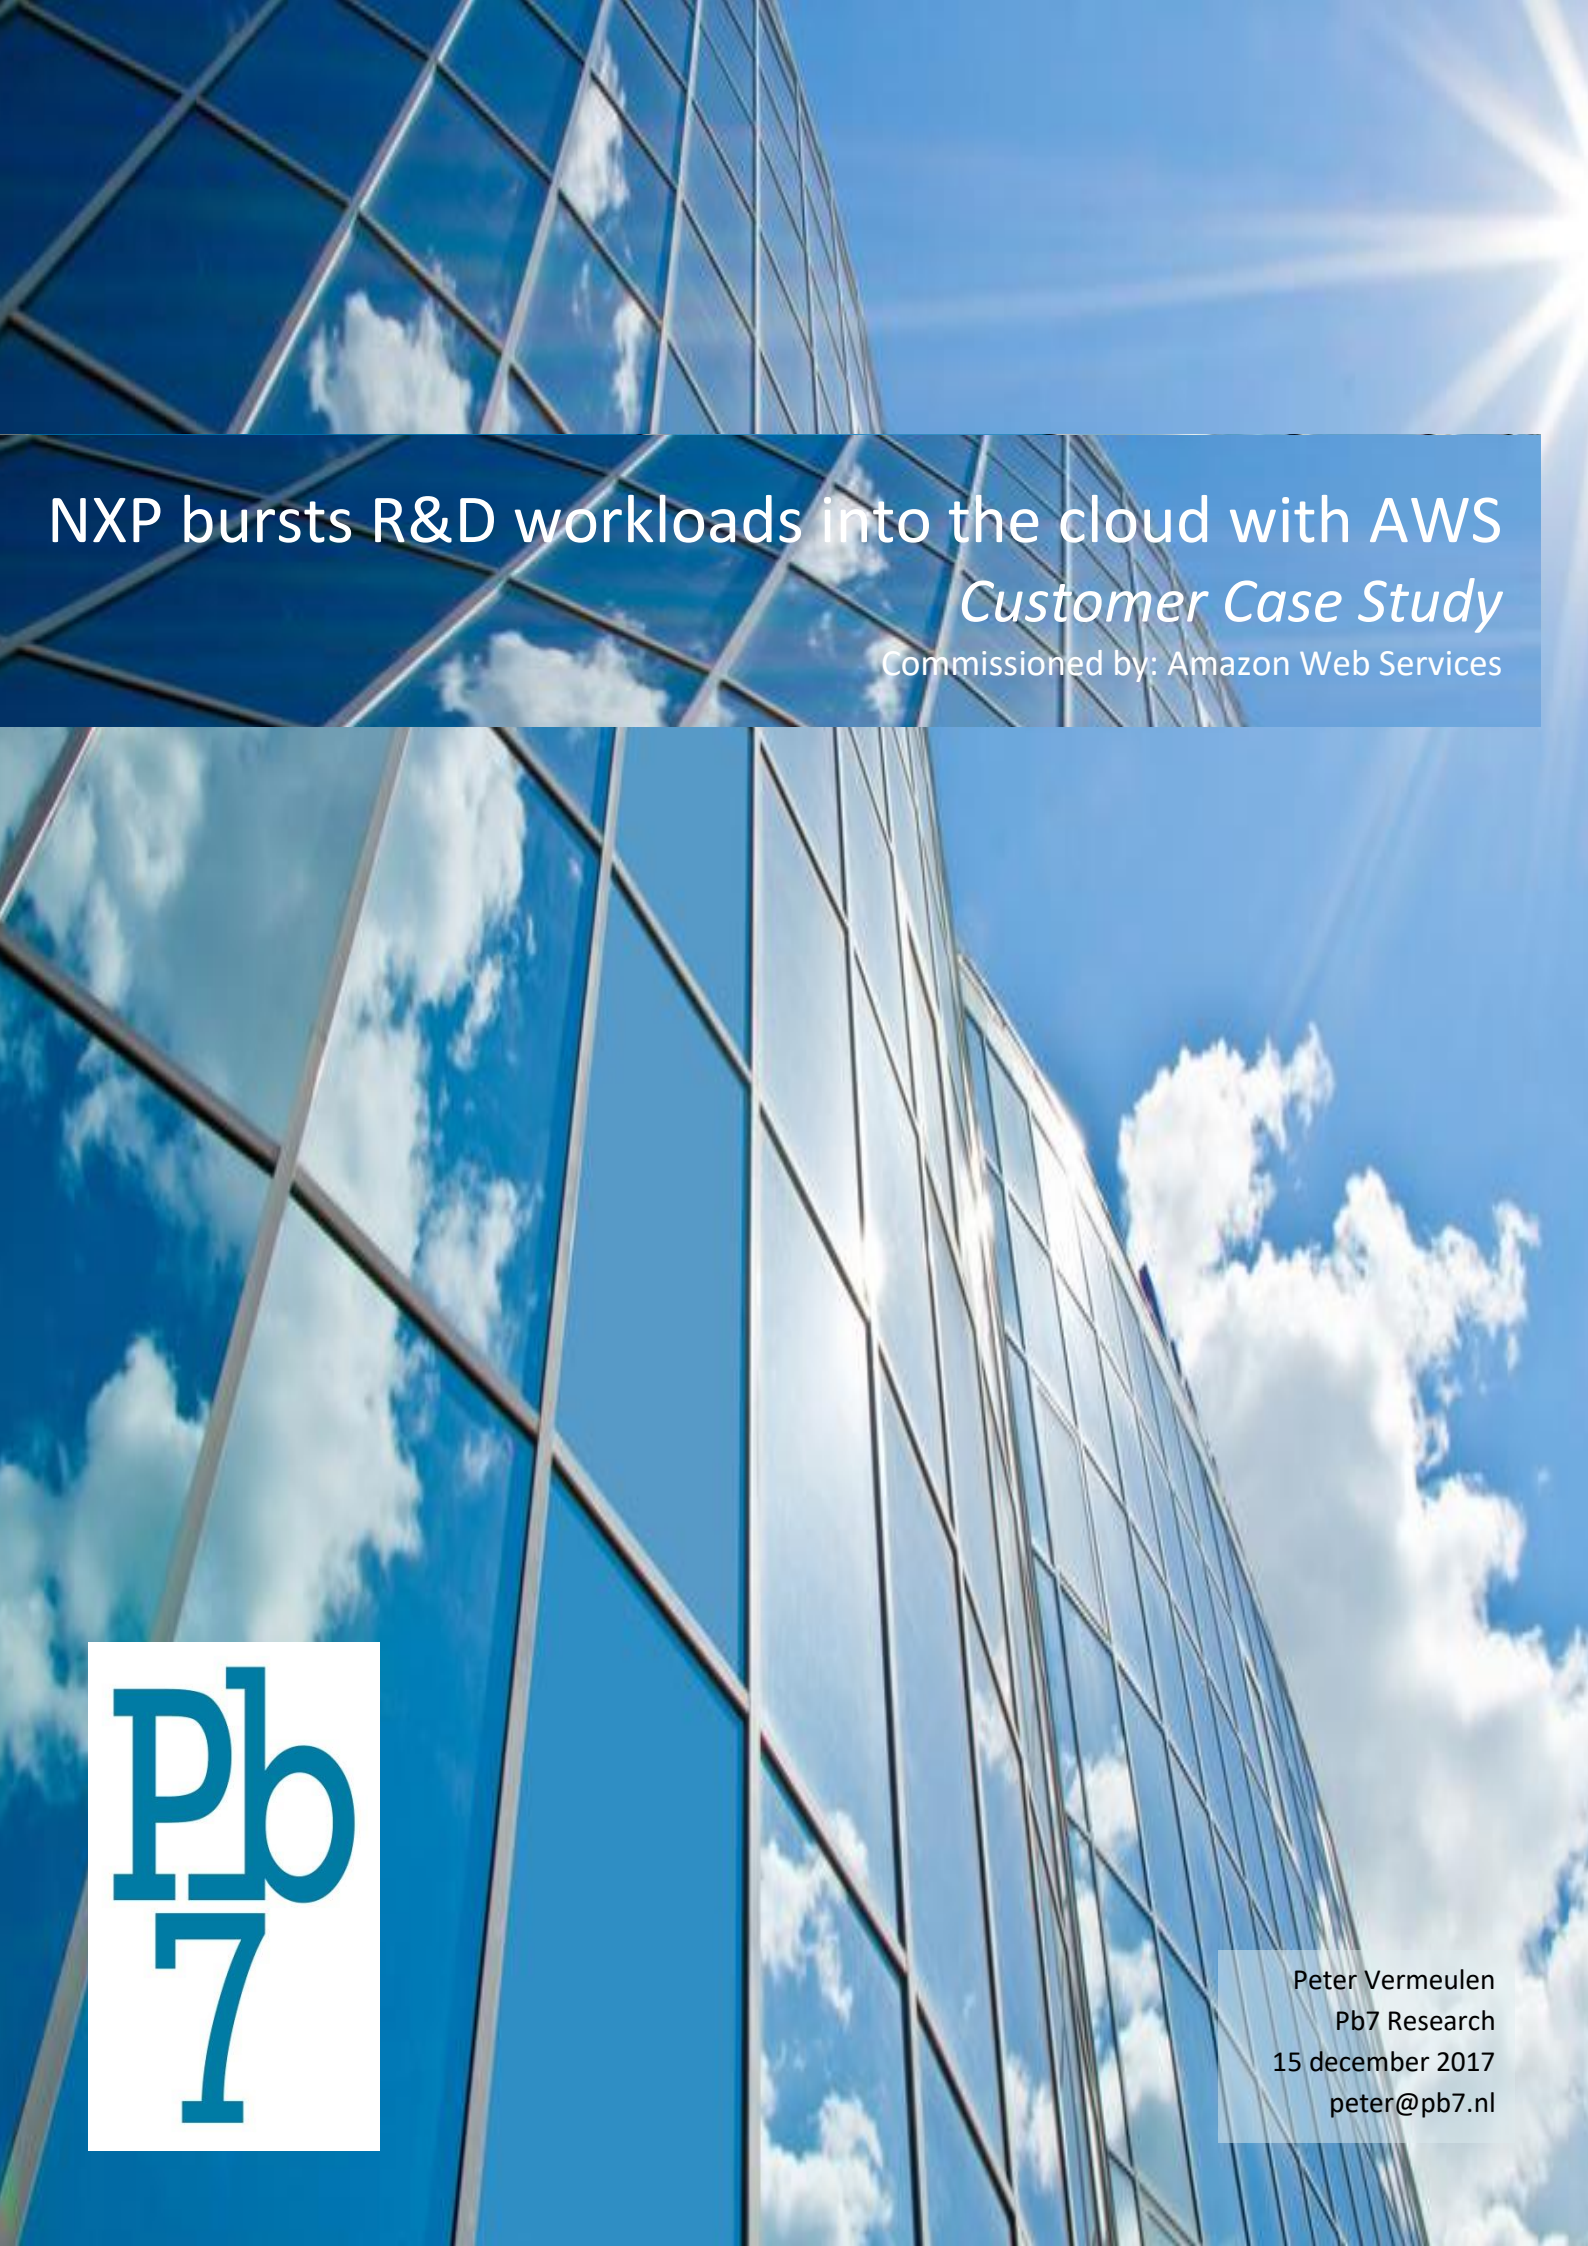

NXP bursts R&D workloads into the cloud with AWS

Customer Case Study

Commissioned by: Amazon Web Services

Peter Vermeulen
Pb7 Research
15 december 2017
peter@pb7.nl

Summary

NXP Semiconductors N.V. is using Amazon Web Services (AWS) to burst crucial R&D workloads into the cloud. In the development process of new chips, NXP needs access to large peaks of computing power on demand. Bursting those peaks into the cloud ensures a short time-to-market and a high quality product. With the help of AWS, NXP achieved those goals, but the road to this hybrid cloud solution has been challenging along the way.

One of the key challenges that led to this solution was the inability or unwillingness of Electronic Design Automation (EDA) software vendors to introduce SaaS solutions based on a pay-per-use model. The lift-and-shift solution was nevertheless very beneficial to NXP. Key to the success were the freedom of the initial team, the gradual introduction of the new platform to the business and "the mainstream" of IT, and the robustness and flexibility of the solution.

For this case study, Pb7 Research interviewed Erik Janssens, Cloud Architect, NXP, and Henk van den Elsen, Program Manager R&D in the cloud, member of Infrastructure IT Solutions and the R&D IT department of NXP.

About NXP Semiconductors N.V.

NXP Semiconductors N.V. is a semiconductor manufacturer, listed on the NASDAQ since 2010 and headquartered in Eindhoven, the Netherlands. It enables secure connections and infrastructure for a smarter world, advancing solutions that make lives easier, better and safer. As the world leader in secure connectivity solutions for embedded applications, NXP is driving innovation in the secure connected vehicle, end-to-end security and privacy and smart connected solutions markets[1].

NXP Semiconductors started out in 1953 as a part of Philips, known as Philips Semiconductors. The company was sold to a consortium of private equity investors in 2006, at which point the company's name was changed to NXP Semiconductors. In 2015, it was announced that NXP Semiconductors would merge with chip designer and manufacturer Freescale Semiconductor. NXP believes that the combination of NXP and Freescale creates an industry powerhouse focused on the high growth opportunities in the Smarter World. The merged entity has leading positions in automotive semiconductor solutions and in general purpose microcontroller (MCU) products.

NXP employs 11.000 engineers across 25 countries and has R&D facilities in the US, Europe and Asia. In total, NXP employs 31.000 employees across 33 countries and posted revenue of $9.5 billion in 2016.

The business challenge

NXP invests heavily in research and development (R&D) in order to provide very high performing, richly-optimized, and very small geometry products. With NXP R&D and the help

[1] The description of the company in this paragraph is provided by NXP

of partners, NXP aims to develop advanced technologies that give customers greater productivity and speed in the time to market, while minimizing the environmental impact.

Simulating a circuit's behavior before actually building it can greatly improve design efficiency and provide insight into the behavior of electronics circuit designs. The compute load for chip simulations is very dynamic by nature and varies on a weekly basis with more than a factor two. But the lead time for any change in the installed compute capacity is vastly slower. This makes it very difficult to optimize the amount of installed compute capacity and deal with peak loads. To optimize the user experience and to decrease pending batches, the R&D IT department started to offload chip simulations jobs in the cloud.

Getting started

At the start of the journey, NXP's IT strategy department identified a corporate need for deploying cloud solutions. A number of workshops were organised to identify the most important use cases. One of these workshops was focused on R&D and it's challenge of the strongly fluctuating demand in capacity. The question that was asked was if it was possible to move the spikes in demand into the cloud. The first thing NXP did was analysing what the current compute usage looked like and where the demand was coming from. NXP made distinctions between types of workload and found that seventy to eighty percent was related to EDA applications. The focus then turned to this type of application.

After NXP R&D IT decided to look for a cloud bursting solution for EDA, they started identifying the best vendor. They first explored moving to a SaaS-application, which was seen as the preferred solution, but independent software vendors (ISV) do not offer a SaaS solution for EDA for semiconductors. In this small and very specialized market, there was – and still is - limited pressure for the ISVs to change their business model and move from traditional on-premises application models to SaaS solutions. As a result, NXP built an IaaS hybrid cloud solution.

During the IaaS provider selection, AWS had by far the widest range of High Performance Compute (HPC) solutions and was the most advanced in terms of functionality. At the start in 2012, R&D IT built a dedicated team, 5 staff who spent about 50% of their time on the project. The team operated separately in an Agile way from the rest of R&D IT and were told to work outside of all existing structures. They started with a credit card, opened up an account, and started testing. The free, agile, and iterative way of development resulted in a high speed of development.

In the beginning NXP started with proof-of-concepts on a small scale. The first working solution was available after 18 months. Because of the complex structure of EDA solutions, NXP did some reverse engineering to find out how the application worked and to identify where in the software stack they could extract the processing and move parts to the cloud. Load sharing software (LSF) was used to move the workload there, and NXP needed to write dedicated software, a data sharing mechanism, to get the data there as well. So as opposed to defining the data requirements in advance, if the application in the cloud asks for data, it will retrieve (only) the required data. This resulted in a secure solution with limited risks.

For NXP that was an important breakthrough in the project, which is basically the essence for a hybrid solution. The security requirements were met by only moving necessary data into the cloud. On the other side, all the tools were moved to the cloud, which led to a number of legal discussions with EDA software vendors. NXP had to explain and commit that the EDA functionality that was built in the cloud would not be available to third parties.

In the following 12 months NXP started to mature the cloud solution for a wider range of workloads and focused on continuous improvement. After 30 months, NXP had a robust solution available. After that, they were confronted with the acquisition of Freescale, that needed to be integrated as well.

Key challenges

The most practical challenge of moving to the cloud was dealing with the EDA software vendors. The ISV was used to determine the license fee based on NXP's data centre capacity. With the cloud in standby, IT had "unlimited supply", but now the department that owned the licenses had a problem, since they needed to get additional licenses. They ran into the problem of the EDA suppliers' license structures that were not based on a pay-per-use model, which is typically used in cloud models.

At the start R&D IT challenged the security department. Henk van den Elsen needed to explain how everything was secured. Van den Elsen argued that NXP had already outsourced its data centres, operations, and support. So there is not really any difference. "We have already accepted that third parties are touching our data. The controls we put in place such as NDAs and screenings can also be applied to cloud partners". And the solution NXP built was essentially an extension of the on-premise data centre. The only thing that can realistically happen is that the NSA will approach AWS with a court order, but they can do the same thing at the door of the on-premises data centre. Even then, all data is encrypted in all states and NXP is the only one with the keys. Further, NXP believes that security in Amazon's data centres is at a higher level compared to NXP's data centres which are already managed by very well-trained, professional service providers.

Lessons learned

Now, five years later, R&D IT is maturing the cloud bursting IaaS solution for EDA. Currently the team is adding additional functionality: making changes to accommodate recent integrations. Also NXP is working on implementing more security controls to allow other type of workloads. The team is very proud of its accomplishments. There were a number of important lessons to be learned:

- *The price of freedom*. Separating the team from the regular IT department was vital to the success. But there was a price to pay at the end of the process as well. When the cloud solution was brought to production, there was a startup mentality, supported by processes and tools. At the end of last year, this DevOps team had to be introduced into the mainstream. Doing that successfully, required spending a lot of time explaining to individual stakeholders what was achieved and how.

- *Hybrid charging*. In the first stages of the project, the project was funded by IT, but the costs increased rapidly and the business would need to start funding it based on a pay-per-use model. Convincing Finance and the business units on that was especially challenging: when you use cloud services, there is a higher need of predictability of business projects. If you start a simulation in the on-premises data centre, it isn't costing any (visible) money, but if you use the (AWS) cloud, capacity needed to be paid on an hourly basis. NXP has now defined a mixed charging model, as NXP IT is acknowledged as a business enabler in providing scalability and faster time to market. At the start of a project, the project team, with the support of R&D IT, predicts how much capacity they need and what that would cost. During the project, R&D IT monitors and reports back how much is being consumed.
- *Cost control*. Keeping cloud costs under control starts with making sure you are not paying for resources you are not using. NXP built an autoscaler to make sure of that. It checks if there is capacity on-premises first. If not, the job will be started in the cloud. Next to that, cloud usage reporting is made available to business units.
- *Redundancy considerations.* At the start, NXP worked with a single region, EU (Ireland). Later, the solution was changed to support multiple regions and US East (Northern Virginia) was added. There have been discussions about creating more redundancy, but in the unlikely event that AWS is not available, NXP can make due with it's internal resources. That could lead to a temporary reduction of capacity and speed, which is acceptable for a limited amount of time. In that sense, the EDA application is not yet business critical.
- *Connectivity considerations.* When NXP started out, they used a VPN tunnel to connect. But as a result of the strong growth over the last couple of years, they decided to implement direct connects with AWS (AWS Direct Connect) in both Europe and the US last year.

Working with AWS

During the process, NXP operated fairly independently from AWS. At a certain point they had an issue with the performance and got in direct touch with the product development team of Amazon, of the EDA provider's team, and Red Hat to identify the issue. This process was driven and facilitated by AWS. AWS was very proactive in solving all issues and started looking into how they could make their solutions workable and optimized for NXP's environment. At the end of the process, AWS started to dedicate resources specifically to high performance computing in the semiconductor industry.

Apart from that, NXP mainly got in touch with AWS to discuss roadmaps: when are certain instances becoming available? There has always been good contact with account management. But in terms of engineering support, NXP was more reluctant to use it since the turnover of AWS employees there was very high. Still, if an issue arises and AWS is contacted about it, they immediately put their shoulders behind it and solve it as well as they can.

The Benefits

NXP started using AWS's cloud solutions in order to get cost-effective access to peaks in compute power. There are now a number of concrete projects that would not have been completed in time and with high quality (due to enough simulations) without it. Van den Elsen mentionted one project specifically, where the stage of test production cycle (when a sample is produced to determine the quality) could be skipped, reducing the time-to-market with three months.

The growth in compute capacity for NXP is 30-50% on an annual basis. NXP IT is much more able to support the business dynamics. For now, cloud and on-premises compute power are growing in parallel at the same speed as the cloud continues to be the solution for peak capacity only.

Looking forward

For the future, NXP is architecting to support more and more services and workloads in the cloud. With the on-premises data centre, NXP is not able to perform non-disruptive maintenance as there is always a project in a critical stage. With a cloud data centre per project, every project can start with a fresh, completely up-to-date environment, while old projects can continue to run in older, stable environments.

NXP also continues to keep its eyes open for SaaS solutions. If a mature SaaS offering for EDA emerges, it will be preferred compared to the current IaaS setup. Still, they are not keeping their hopes up and don't see that emerging in the foreseeable future.

Conclusion

NXP semiconductors uses AWS mostly as a toolbox for bursting EDA work loads into the cloud. They found that AWS is always willing to commit and proactively solve any issues that arise. For NXP it is important that AWS is leading in terms of functionality and services, even though it is difficult to always keep up as a customer. The biggest challenges for NXP were not in working with AWS, but had more to do with ISVs that are holding on to old business models and with internal organisational challenges.

NXP semiconductors uses AWS to make a real competitive difference. The EDA cloud bursting solution that has been built over the last five years decreases the time-to-market of new chips as it ensures that all design teams have access to plenty of capacity for simulations when they need it. It will be interesting to watch how NXP's vision for future cloud usage materializes.