„NetApp Private Storage for Cloud“ – an opinion on the application of European data protection standards

by certified IT law Attorney Dr. Jens Bücking*

The „legally admissible“ provision of company-critical information is an absolute precondition of „IT compliance“ - regardless of its specific legal basis - with the guarantees of availability as well as backup and archiving processes being crucial to compliance with the applicable standards.

e|s|b Rechtsanwälte

Dresden Stuttgart Berlin Leipzig Prag

NetApp®

I. Executive Summary

The transfer of data to the Cloud always requires legal legitimation. Especially in relation to transfers to non-secure „third-party states“, as the USA is classed by the EU in matters of data protection, most providers have previously primarily relied on „Safe Harbor“ self-certification. Now that the ECJ has heralded the end of „Safe Harbor“ in its ruling on 10/6/2015, the most prominent solution to this data protection issue is no longer available. Even the use of EU standard contractual clauses has been cast into doubt by the ECJ‘s ruling. These two principles, however, formerly provided the basis permitting many well-known Cloud providers to transfer data to the USA. The only remaining option to avoid this legal uncertainty and to ensure an appropriate level of protection which is simple to implement is the hybrid Cloud solution, as offered by NetApp in its „NetApp Private Storage for Cloud“. This solution gives the company, as „master of its data“, a simple and legally admissible way to continue benefiting from the advantages of the Cloud without any concerns. The data are stored in the relevant country according to the relevant requirements and the compute service can be accessed from the public cloud.

II. Threat scenarios

The issues of data security and availability are now an inalienable part of corporate reality. While the protection of intellectual property and innovation used to be the decisive factors in market endurance and success, in an age when data and systems are networked globally the constant availability of information (data) has taken over as the most important factor. International studies show that 70% of companies which suffered catastrophic losses of data were forced to shut their doors within 18 months. In the knowledge that the availability of corporate data is a „do-or-die“ criterion, numerous national and international (special) laws and regulations, e.g. various national data protection laws and laws regarding the rights of corporations such as the Sarbanes Oxley Act (SoA or „SOX“) in the USA, the „Euro-SOX“ statutory auditor regulation at the EU level and „Basel/Solvency II et seq.“ at the international level, now oblige companies and - under comparable regulations - public institutions to implement effective risk and information management. Compliance with these is now a management issue and, as part of „corporate governance“, is one of the most important corporate leadership tasks. This raises questions of personal liability on the part of the management, compliance and security officers and other decision makers in key positions. In the event of significant breaches of obligations, this can also affect the insurance coverage of the company and members of management.

It therefore seems that contracting data and system maintenance to certified specialist companies and/or outsourcing certain types of processing to globally accessible computing centers would be sensible for more than merely economic reasons. From the point of view of risk, however, on the basis of recent frightening statistics about organized business espionage this would only seem justified in certain very limited circumstances - i.e. in secured EU/EEA locations with appropriate security certificates and the relevant information and audit rights relating to cloud providers, etc. Numerous companies were, however, still recently using the possibility of a „Safe Harbor“ certification to justify the transmission of data to the USA. This is because the transfer of data to the USA (considered an insecure „third-party country“ in data protection terms) was - regardless of the need for a legal basis - only permitted when the recipient was able to provide a suitable level of data protection and confirm this by means of the corresponding certification. According to the (de facto decisive) statement of the „Article 29 data protection group“ (as independent advisory committee to the European Commission on issues of data protection), however, there were significant doubts as to the effectiveness of „Safe Harbor“, since the current implementation of this certification was seen as not (or no longer) suitable to provide an adequate level of data protection - and that „Safe Harbor“ should no longer provide „cloudsourcing“ companies and their decision makers with carte blanche in terms of liability. When various national data protection authorities (such as the German „Düsseldorfer Kreis“ as the leading data protection committee for the federal and state governments) subsequently criticized the „Safe Harbor“ principle as insufficient, the European Court of Justice (ECJ) ruled on 10/06/2015 (C-362/14) that, in practice, this regulation

was no longer valid for the exchange of data between the USA and the EU. It stated that the USA, as a „third-party country“ without a level of data protection corresponding to that in the EU, failed to provide sufficient protection for personal data even with „Safe Harbor“. 4,410 US companies are currently „Safe Harbor“ certified, including major cloud providers such as Microsoft, Apple, Adobe and Google.

In the event of the transmission of data to the USA, the entity responsible for processing the data must, effective immediately, no longer rely on self-certification in accordance with „Safe Harbor“ principles and instead verify each individual data transmission‘s compliance with legitimacy criteria, make use of the EU standard contractual clauses relating to data processing or comprehensively (i.e. group-wide) subject itself to binding corporate rules (BCR) for data protection and have these approved by the responsible data protection authority. What remains to be seen is whether the EU standard contractual terms will measure up to the strict requirements which the ECJ also established in its ruling, since supervisory authorities may, in individual cases, forbid the data transfers on this basis if they consider that the rights of the persons affected are not granted sufficient protection.

All of this seems to preclude any data protection-compliant outsourcing of the processing of personal data to the market leading clouds, which generally rely on „Safe Harbor“ and the EU standard contractual clauses and which are predominantly US companies. The data management solution market has, however, proven itself to be innovative. The focus now is on hybrid and multi-cloud architectures which seek to provide a suitable answer to the open question of a compliant global cloud application. What follows is a look at whether, and under which conditions, this question can be answered:

III. Current trends and challenges in IT

The assessment of data for the analysis and implementation of business processes is currently one of the greatest challenges facing IT. Due to growing quantities of data and higher required processing and transmission speeds, conventional enterprise processes can no longer be effectively managed solely „on premises“ with a reasonable use of infrastructural resources. Company data must therefore be stored on a system that is able to combine high availability with a high level of security against failure and attack - all under the continuous control of the authorized and legally responsible company as „master of the data“. This means that the use of cloud infrastructures is now unavoidable for the provision of contemporary IT environments.

Modern strategies must consider the use of multi-cloud environments - even though cloud providers may vary greatly in terms of price, performance and range of services. Digital business processes now require maximum flexibility and scalability in terms of storage capacities and processing speeds. At the same time, the entire system must be „hardened“ as the threat of cyberattacks increases exponentially. This leads to the next condition which arises from the legal framework conditions. The allocation of tasks to multiple cloud providers may not result in a reduced ability to monitor the data to be processed. Companies must be able to ensure that they are not only aware of the provider and subcontracted service provider, locations and manner of the data processing at all times, but also that they maintain control at all times. Under the perspective of legal liability, this applies both to personal data as well as to other business-critical information such as intellectual property, research/development data, evidentiary correspondence etc.

These challenges can no longer be effectively met with conventional cloud computing and traditional IT architectures. In addition to fulfilling the legal framework conditions, management must ensure technical and organizational protective measures as well as a flexible, scalable and always-available high-performance system, all while maintaining continuous - and documented - control at all times. The problem is that in the conventional outsourcing model the data are stored and processed in the cloud and not out of the cloud – which, as will be shown, is where the path leads for the transmission of certain categories of data in clouds outside of the EU/EEA in light of

strict legal framework. If data are transferred to the cloud, however, companies often face the problem that a later restore can no longer be guaranteed. The same applies to the vital ability to delete data stored in the cloud which - despite assurances from the cloud provider – is de facto frequently not possible.

This clearly marks the path towards hybrid and multi-cloud architectures in which specific categories of data are kept available on a dedicated own storage system located in close proximity to a network node at a co-locator‘s premises and remain under the company‘s complete control. In this case the company remains master of its systems and data, which it operates at the co-locator‘s premises while maintaining continuous sole availability and constant functional and process control. The data processing itself therefore does not take place in the cloud, but rather out of the cloud, which from a data protection perspective gives reason to reassess the conventional models of data transmission and contracted data processing and - as will be shown below - can lead to significantly more leeway for the use of global clouds.

In terms of the personal nature of many categories of data, the actual processing is, of course, limited to „blind“ computer operations on the temporary mass storage of the cloud provider‘s virtual machine. The volatility of the storage can therefore guarantee the deletion of all data as soon as the application initiated by the company is completed, the virtual machine is shut down or suffers a technical failure. What is needed is therefore a cloud structure suitable for use with existing enterprise applications. The mobility of data, in particular the ability to restore them, must be guaranteed - with continuous control by the company - throughout the entire duration of the processing from cradle to grave. That is to say from the first „computing job“ out of the cloud to the return transmission and „residue-free“ deletion from the cloud provider‘s systems, with corresponding guarantees from the provider. If these conditions can be fulfilled and technically and organizationally guaranteed by the provider (and ideally insured as well), such cloud infrastructures can provide a real alternative to on-premise structures from the perspectives of both economy and performance.

IV. NetApp Private Storage for Cloud

Hybrid cloud solutions typically combine a private cloud with public cloud environment, thereby temporarily offsetting or permanently expanding the technical limitations of an on-premise IT infrastructure in terms of performance and storage space into the cloud by means of external resources and services from various providers in a multi-cloud model.

With „NetApp Private Storage for Cloud“ and its „Cloud ONTAP“ development, one of the leading providers of data management solutions now offers the kind of hybrid multi-cloud described above out of which cloud resources can be used - while simultaneously considering requirements relating to the mobility, availability, security, confidentiality and continuous control of the application-relevant data. Cloud ONTAP is a storage operating system which can be purchased „via mouse click“ in the Amazon Marketplace and installed and made available as a virtual appliance on an Amazon EC2 instance (virtual machine). The client then receives a link to the configuration screen. Cloud ONTAP then works as a storage controller for data management within a virtual environment in the Amazon cloud which is created and operated there for the relevant companies. Billing takes place by the minute on the basis of a consumption model. All data are managed by Cloud ONTAP and stored in encrypted form on Amazon Elastic Blog Storage (EBS) volumes.

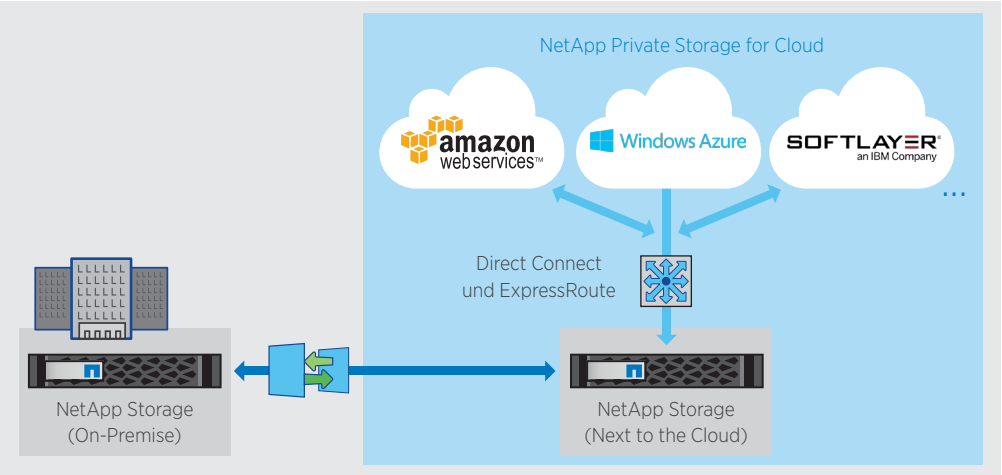

Figure 1: NetApp Private Storage for Cloud enables a Hybrid-Cloud-Solution, with which companies can store their data close to the cloud and use cloud resources while maintaining complete control and safety.

The solutions offered by NetApp also include the (private) „Cloud on the Cloud“. If a „Clustered Data Ontap“ is used within the on-premise enterprise infrastructure, this permits certain services to be expanded into the cloud. A storage system under the exclusive authority of the client is installed in a co-location computing center to store the company data. The computing center is in direct geographical proximity to the provider, ensuring the lowest possible latency and therefore higher access speeds at network level. Data mirroring is set up between the on-premise NetApp system and the system installed at the co-locator‘s premises so that the data from the co-location system are made available on the cloud for processing – with correspondingly low latency times. Direct IP connections are established between NetApp Private Storage for Cloud and the individual public cloud providers. These are exclusive data connections that do not run via the Internet. The public cloud provider‘s services and resources can then directly access the NetApp Private Storage client‘s data. The processing takes place without the data being pushed to the provider‘s public cloud.

Such innovative solutions mean that it is no longer necessary to move the data. Instead, these remain safely on the storage system which remains under the unlimited and continuous control of the company. The reserved public cloud services (computing power, applications and other „as is“ services) then access these data to carry out specific types of processing at defined times - as determined by the company itself. At the end of the data processing the services then log out or - metaphorically - pull out of the data sets which they had drawn into their working memory for processing. The processing of course made changes to the data and used them to produce results, carry out evaluations etc. At no point, however, was there any „migration“ - or a „data transmission“ in the legal sense - into the cloud, since the storage location of the data specifically did not shift towards the public cloud for permanent storage and instead - quite the opposite - the company drew the cloud services to itself out of the cloud. Only the computing process itself, as opposed to any storage space, is accessed in the cloud. These are merely computing jobs which can be outsourced via specific processes. The data themselves naturally remain on servers owned by, or (in ownership models such as rental) in any case under sole power of disposal, similar to ownership, of the company.

V. Data transmission/Contracted data processing vs. Infrastructure expansion

On this basis, no transmission of data in the legal sense actually occurs and - in a step which currently sets the NetApp solution apart from other international cloud models - there is not even a single instance of contracted data processing.

After all, the process merely involves the „extended arm“ of the on-premise enterprise processes: an external application that only „starts up“ when instructed to do so by the owner (and/or sole authorized user) and „master of his data“ and, therefore, the party „responsible for the processing“ in European terms. As such an „extended arm“ of the in-house IT infrastructure resources which „reaches out to the cloud“, there is naturally no transfer of data to the third-party responsibility of the cloud provider. The crucial distinction is between „store“ and „compute“: The provider does not rent out any storage, merely computing power, and its systems merely carry out „blind“ data processing. The conventional cloud-typical (return) transmission of the data to be processed to the provided cloud structures does not take place. This process, which controls the actual data processing on its own systems, does appear more fleeting as it takes place solely in the working memory. The computing is brought in-house and not „cloudsourced“.

This difference is decisive. Contracted data processing required - in this respect barely practicable - preliminary examinations as well as the conclusion of extensive contracts with the various cloud providers. This obligation applies under the European Data Protection Directive (96/46/EC), which describes minimum data protection standards to be implemented in EU member states by means of national laws, as well as its upcoming legal successor, the EU Basic Regulation on Data Protection (DS-GVO). These state that contracted data processing is only permitted under European law if a law allows it, or if a written contract regulates the contracted data processing in detail.

If no such contract exists, this represents a breach of the client‘s obligation to document the respective responsibilities in written form and may lead to penalties. The contracted processor must be carefully selected by the party „responsible for the processing“ according to criteria such as reliability, performance, technical and organizational security measures which correspond to the current state-of-the-art, etc. The selection in particular entails verification of the technical and organizational measures prior to the beginning of the data processing as well as throughout its duration. The relevant contracted data processing contract must fulfill certain minimum requirements in terms of content. It must in particular establish that the contracted processor and any persons subordinate to him with access to personal data may only process said data at the instruction of the „party responsible for the processing“. Regardless, the „party responsible for the processing“ must remain able to legally and factually intervene in the decisions of the contracted processor regarding the means of processing; overall responsibility remains his. Logically, it follows that the „party responsible for the processing“ must effectively monitor all contracted processors in order to ensure that their decisions comply with his instructions, the contracted data processing and data protection laws. Should a contracted processor fail to comply with the limitations of the stated use of the data, he then „mutates“ to become an - illegally acting - „party responsible for the processing“. The original „party responsible for the processing“ thereupon must explain why and how the contracted processor breached the conditions of the contract. In such cases the Article 29 Data Protection Group tends to favor joint legal responsibility (and liability) as this achieves the best possible protection for the interests of all those involved. One important consequence of joint responsibility is the joint liability for damages. To this is added the fact that, while the EU Data Protection Directive primarily holds the „party responsible for the processing“ liable, other than in the exceptions described this does not prevent individual national data protection laws from also making the contracted processor liable in certain cases.

Now that „Safe Harbor“ can no longer be used as legitimation and the use of US clouds is now fundamentally considered to entail the transmission of data to a non-secure third-party state, according to the Article 29 Group what is now needed is an additional agreement regarding contracted data processing with the aforementioned content regarding the fulfillment of carefully considered requirements relating to the need and appropriateness of such a data transmission. This would, however, impose almost insurmountable obstacles on any permissible data processing in US clouds.

There are, however, many arguments in support of the idea that the hybrid cloud solution described, as made possible here by „NetApp Private Storage for Cloud“, can be seen as neither a data transmission nor contracted data processing, but rather as an extension of a company‘s IT resources (and therefore as the extended arm of its own infrastructure and data processing systems):

The „collection“ of the data to be processed by cloud services or applications stored in the cloud from a co-located client machine does not represent a data transmission in the legal sense as interpreted by the courts and data protection authorities. This is because the crucial term „transmission“ should not be taken to mean a „transport“ to a third party, i.e. an export of the data. A transmission should rather be understood to mean a „disclosure“. What matters is the distribution of the information contained in the data. The manner in which the data are presented or accessed and the details of the information transfer are unimportant because each increase in the size of the group which has access to it touches upon the interests of those affected. This includes any action which results in the data arriving in the addressee‘s area, regardless of the actual details, meaning that a call procedure is also included. A data transmission would also be the case if access to the data is transferred and the data are therefore usable by the recipient. In this case the transferring party must keep the data available for this purpose, that is to say have created the necessary technical and organizational arrangement to do so, but the decisive activity is undertaken by the recipient. The data transmission is the viewing and/or access, that is to say the actual taking of possession of the data as a result of the recipient‘s action.

Naturally, no viewing and/or access in the legal sense takes place in the absence of the notification of the third party because the data processing takes place as if in an ephemeral „black box“ without third-party access and usage rights - as the aforementioned extension of the company‘s IT resources. This pure infrastructure use can therefore not be considered a data transmission. And, of course, it does not even represent contracted data processing if the party responsible in a legal sense - i.e. the company that must ensure control of the data at all times in order to control the purpose and means of processing - neither transfers this responsibility nor expands it by integrating a contracted data processor who must be trained and instructed, but rather makes use of third party infrastructure. This is the case if data processing systems are merely rented or otherwise used and these are operated by the responsible party and the „landlord“ - in this case Amazon - would only be able to access the data by violating rights of ownership or disposal. The important factor here is the functional control of the data processing as decisive criterion for the control of technical systems.

This functional approach, which the Article 29 Data Protection Group also supports in principle, can also be applied to cases in which computing power is rented on external servers and the responsible party then personally uses it to process personal data. If a computing center wholly or partially places its systems or services at the disposal of a client for individual data processing activities and the latter uses them online to carry out data processing shielded from third-party access, what has taken place is not contracted data processing, but rather independent data processing by the client. This must in any case be assumed if, as in the example of NetApp Private Storage for Cloud, the client solely and exclusively decides which, when and how data are processed and also determines the necessary programs and algorithms. The cloud computing center is limited to ensuring readiness and keeping track of the duration of use. In these cases there is no use of contracted data processing. Responsibility for data protection remains with the companies.

In the view of certain European data protection experts, contracted data processing does take place in cases in which the contractor merely makes IT resources such as storage or servers available which the client can then use via the Internet since, in these cases, as a rule the contractor would have the technical opportunity to affect and access the personal data in question. The Article 29 Data Protection Group of the EU Commission has not yet issued an opinion on this. If the approach is followed to its logical conclusion, however, nothing would change about the conclusion that the arrangement described above is not contracted data processing. This is because, unlike in the general case described, the hybrid-cloud approach does not entail any such opportunities to affect and access the data. If, and to the extent that, the service provider has no access to the data being processed on a particular storage medium (similar to server housing), according to this approach all that is occurring is an extension of the company‘s responsibilities (as „party responsible for the processing“) to include additional infrastructure resources, and not any type of contracted data processing.

VI. Summary

Solutions such as NetApp Private Storage for Cloud, which was developed in collaboration with Amazon, Microsoft, IBM and various co-location providers, provide a hybrid cloud storage infrastructure which maintains the same level of control as on-premise architectures with the same advantages as public cloud services. The processing takes place without the data being pushed to the provider's public cloud. This means that the migration of data between different providers and data synchronization are no longer required. Since the data remain stored on the company's own storage system, complete control remains guaranteed. As a technical extension of the IT resources under the company's sole control, this arrangement is not subject to the strict requirements for permissible data transmissions to third-party states and does not entail contracted data processing which must be comprehensively contractually and technically secured.

* The author is an attorney and specialist in IT law. He is also the founder of law firm e/s/b Rechtsanwälte (http://www.kanzlei.de), a specialist author on IT law, lecturer at the University for Applied Sciences in Stuttgart and associate professor at the E.N.U. in Kerkrade, the Netherlands.