



How to resist Data Gravity

Ensuring Data Control and Global Scalability with
Hybrid Cloud and Multi-Cloud Architecture

A Strategy Paper created on behalf of NetApp Deutschland GmbH

By Rene Buest, Senior Analyst & Cloud Practice Lead

Executive Summary

- As a result of growing digitization of business models and processes, CIOs and IT decision-makers are compelled to seek new sourcing models and infrastructure concepts. The Public Cloud plays a major role in this context.
- In the digital age, the significance of data management takes on a new dimension. Data has certain inertia and, hence need to be assigned to different classes so that they comply with regulations, legal requirements, technical limitations and individual safety classes. This so-called „Data Gravity“ impacts the mobility of data.
- New storage concepts are needed to process these hard-to-move data outside of one's own IT infrastructure without loss of control.
- Hybrid and multi-cloud storage architectures provide implementation strategies and robust usage scenarios that are aligned with the new requirements. Within these architecture concepts, the data is located in a company-controlled area and the data owner is the only person who determines which parts are to be stored in the public cloud.
- Accordingly, all benefits of public cloud infrastructures can be utilized without losing control of one's data, while fulfilling the necessary compliance guidelines and legal requirements.

Table of Contents

1. Enterprise IT in the Digital Age	4
2. Hybrid Cloud Architectures & Data Gravity	6
3. Enterprise Storage in the Multi-Cloud World	8
4. Hybrid Storage Concepts in Business Practice	10
5. Outlook	11
About NetApp	12
Author	13
About Crisp Research	14

1. Enterprise IT in the Digital Age

Digital transformation tops the agenda of today's CIOs. The focus here is on the reassessment of existing business models as well as on the modification and redesign of entire process chains. In this context, IT organizations not only act as enablers and important initiators but also need to make sure that the potential of new technologies is being exploited. The goal is to create conditions that foster change in the existing business models and enable new digital business ideas. For this reason, business processes and workflows need to be transferred into the digital world.

As part of the digital transformation, IT organizations are also subject to a profound transition. One of the central building blocks of this change is the realignment of the setup and operations of IT infrastructure. Here CIOs face different challenges, such as:

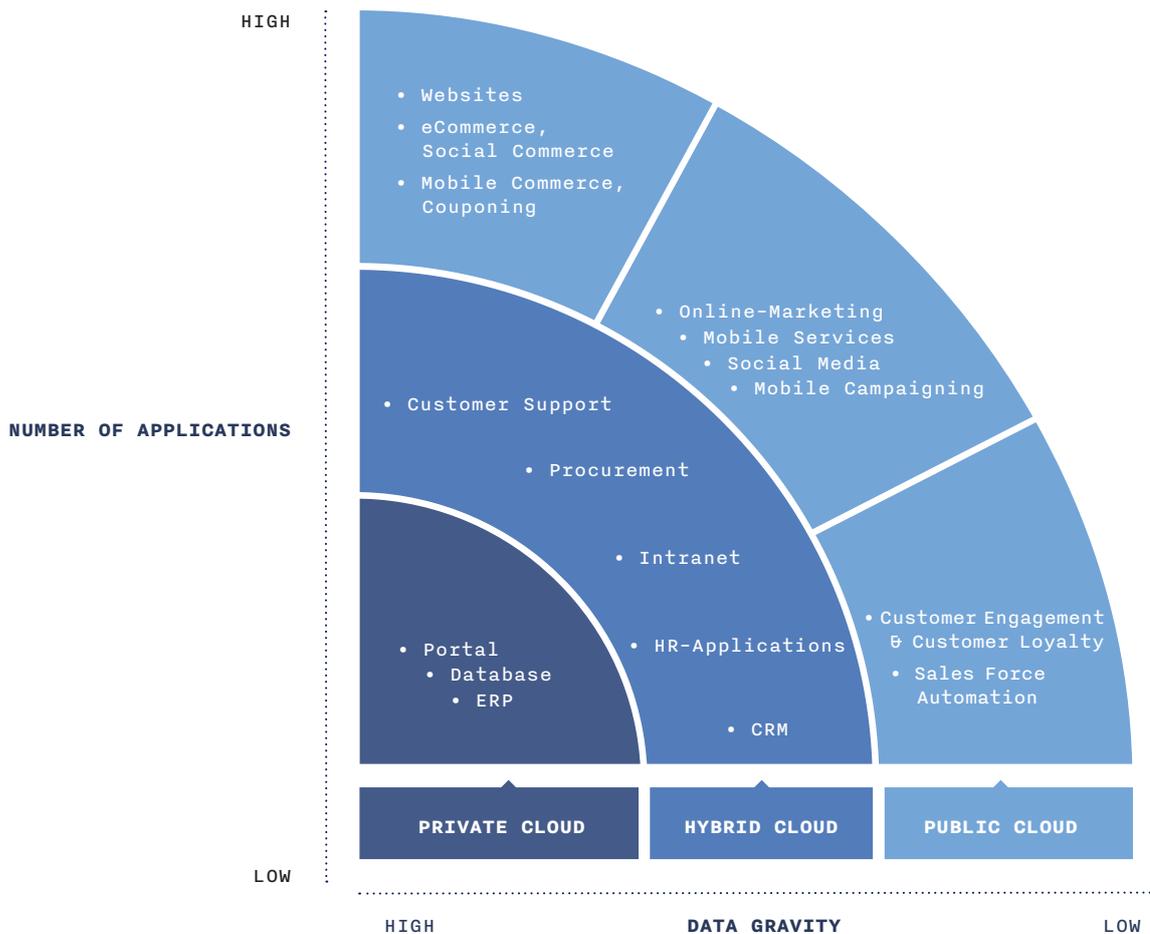
- **Complexity:** The use of cloud infrastructure is absolutely essential to the provisioning of modern IT environments. Only the cloud makes it possible to efficiently provide IT services on demand in a self-service mode to employees, business partners and customers. The challenge here is that the design of the relevant infrastructure is sophisticated and eventual cloud services are based on a brand new, "cloud ready" applications architecture. On top of this, in both cases specialized knowledge, which presently is a scarce market resource, becomes a necessity.
- **Data silos:** Only rarely is a single cloud service being used. Instead, cloud strategies take into consideration the use of multi-cloud environments. There is a good reason for this - price, performance and above all the portfolio are ultimately the distinguishing features of providers. Only a well-balanced provider choice ensures a perfectly tailored cloud environment. Inherent to this perfect cloud mix, however, is the massive challenge of managing data in a way that no data silos are created, because eventually the networking of data is designed to generate new business advantages rather than "forgotten" assets.
- **Scalability:** Digital business models and advanced data-hungry services demand continuous capacity increase with regard to performance as well as on the storage level. The infrastructure needs to be able to dynamically compensate varying loads, if required, while also quickly exploiting other geographical markets (global scalability). Furthermore, the storage infrastructure always needs to provide sufficient memory to the applications and workloads.
- **Security:** The growing degree of digitization and, consequently, the intensified interconnection of all systems, cause the threat from cyber attacks to grow. Not only external attackers play a pivotal role here - so-called insider attacks are gradually moving into focus, too.
- **Control:** The vast degree of interconnection leads to a weakening of the control effectiveness. One of the reasons is the use of distributed infrastructures, which are no longer within the immediate area of influence. Another factor is the growing degree of freedom that employees have to store their data using countless cloud services. This brings about challenges with respect to compliance and meeting legal requirements such as data protection. If it is also desired to have data available in other countries and to process these data, basic legal issues have to be clarified. Customers may be required to prove where their data are located (for example, bank data or health insurance data). Apart from this, the protection of critical assets like intellectual property and R&D data is ensured.

These challenges cannot be addressed through conventional methods and IT architecture. New approaches are imperative to ensure that the legal framework as well as all necessary protective measures are in place, without limiting the productivity of the business at the same time.

2. Hybrid Cloud Architectures & Data Gravity

In the digital age, the significance of data management takes on a new dimension. On the one hand, a growing number of analyzable data are generated and collected. On the other hand, the vulnerability of data increases because of potentially high interconnection of partly unknown systems and distribution of data to cloud infrastructure that is no longer under one's own control.

Data Gravity in the Cloud



QUELLE:
Crisp Research AG, 2015

Data Gravity

Depending on characteristics and security class, data have specific inertia. This so-called “Data Gravity” causes differences in data mobility properties. The main reasons for data gravity are:

- Data size
- Latency of data access
- Throughput/ Bandwidth for the data access
- Controllability of the data
- Fulfillment of legal requirements
- Access to local / mobile apps and cloud services

Hybrid Cloud Architectures within the Framework of Data Gravity

Hybrid cloud architecture usually connects private cloud infrastructure to one or more public cloud environments. Hence, companies can temporarily overcome the limitations that their own infrastructure imposes with respect to computing power and memory, or permanently expand their infrastructure by means of IT resources and services from different cloud providers in a multi-cloud model.

Within the framework of “Data Gravity,” hybrid cloud architecture takes on an important role, as this architecture type is able to address the challenges of data gravity and to provide distinct solution scenarios. In this context, it is no longer necessary to move the data. Instead, it remains in a storage system under one’s control. Public cloud services in the form of computing power, applications and other services access these data during runtime. The services cease when processing is completed. Although alterations to the data are made during processing and results on the basis of these alterations are drawn, the data never leave the storage location to migrate to the public cloud.

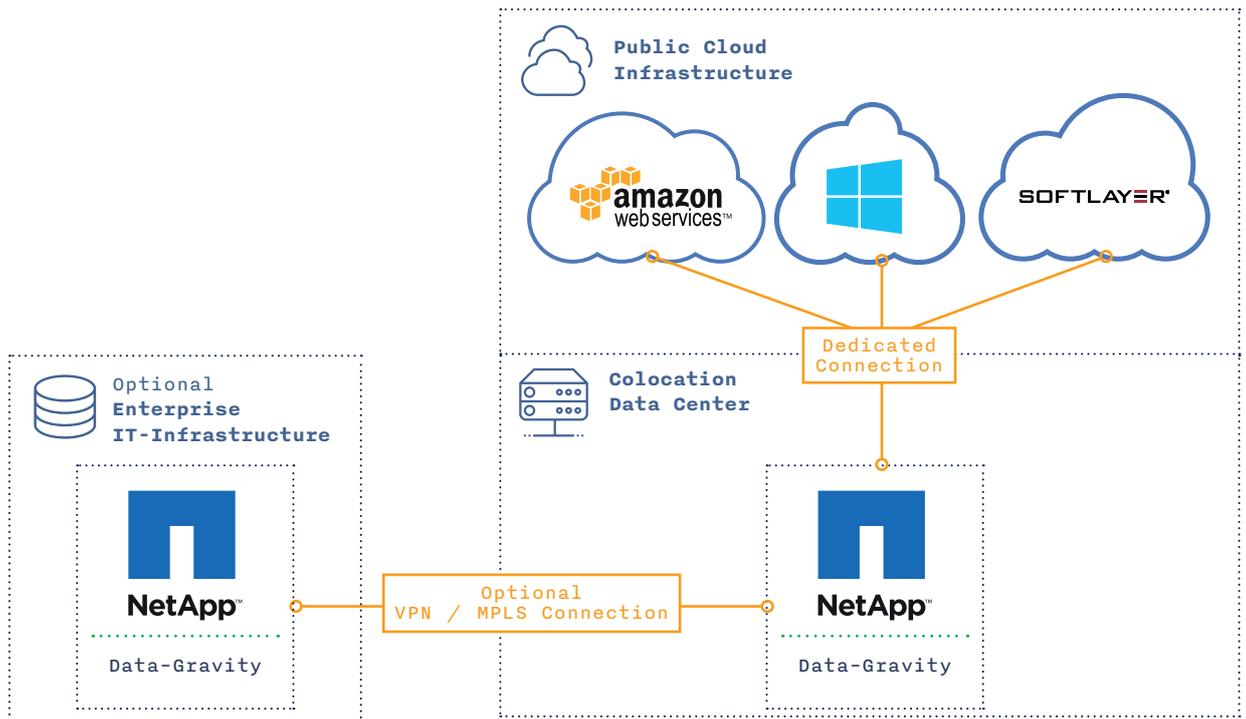
This means that a hybrid cloud scenario of this kind is an ideal basis to resolve current challenges during data processing in the cloud through providing:

- Performance stabilization or improvement
- Reduction of latency
- Increase in availability
- Access control optimization
- Adherence to compliance regulations & legal requirements
- Observance of governance

3. Enterprise Storage in the Multi-Cloud World

Public cloud infrastructure opens up various ways for companies to develop innovatively and to benefit from their data assets. However, compliance and governance regulations as well as the required control level turn out to be obstacles in many cases. The architecture approach of the “NetApp Private Storage for Cloud” stands out as an innovative concept in overriding these central elements of “Data Gravity”.

Hybrid Cloud Storage Scenario NetApp Private Storage



QUELLE:
Crisp Research AG, 2015

In cooperation with Amazon Web Services, Microsoft Azure, IBM SoftLayer and co-location providers like Equinix or e-shelter, NetApp has developed a hybrid cloud storage scenario where the control level is kept as strict as possible, yet allowing for the advantages of the public cloud to be used to their full extent. For this purpose, a customer-owned NetApp storage system that contains company’s data is placed in a designated cage in a co-location data center. The co-location provider uses a cloud hub and secures partnerships with public cloud providers; the data centers are also located near the providers. In this way, low latency is ensured, which in turn enables high access rates on the network level. A direct connection is established between the “NetApp Private Storage for Cloud” and the individual public cloud providers. This connection consists of exclusive data links, which are not directed over the Internet. Thus, services and resources of the public cloud providers have direct access to the enterprise storage data.

The data is processed without transferring them to the public cloud of the provider, staying in the company’s domain. The “cloud hub“ model follows a provider-neutral approach, which is why multi- cloud scenarios can be mapped out. This setup allows for new public cloud providers to be added quickly and connected to the enterprise storage.

Furthermore, on-premise enterprise storage systems can be connected with a designated NetApp storage within a co-location data center via an MPLS link. Both systems can be synchronized continuously by means of this link, thus ensuring less data traffic and helping to reduce costs. The latency for the data access can be reduced as well, depending on the location of one's own on-premise infrastructure.

The restriction of data gravity can be overcome with the "NetApp Private Storage for Cloud" approach. Some of the most important advantages are:

- Multi-cloud access with low latency and high performance. The services of the cloud providers access the data without moving them so that data migration between the providers as well as data synchronization are not required.

- The data is placed in the company's own storage system. This allows for unrestricted authority of control and has a positive impact on the observance of compliance regulations.

- Combining the benefits of using the public cloud with the advantages of enterprise storage.

4. Hybrid Storage Concepts in Business Practice

Hybrid storage concepts are suitable for a large number of application cases. The following examples are anonymous but nonetheless real application scenarios, which have been implemented using the “NetApp Private Storage for Cloud”.

Industry	Scenario	Scenario	Solution
Manufacturing Industry	SAP Disaster Recovery	The SAP infrastructure of a medium-sized manufacturing company is to be operated in a fail-safe manner. A disaster recovery concept was chosen for the implementation of this goal. This means that in the event of a failure of the company's own IT infrastructure, the existing SAP environment continues to operate seamlessly.	First, the SAP data was replicated to a NetApp storage in a co-location data center. In this way, the data remained under the company's control. Virtual machines with SAP applications were then preconfigured but not started at Amazon Web Services (AWS). If the SAP infrastructure at the company location switches into a fault condition, the AWS instances are started up. These instances access the previously replicated SAP data in the NetApp storage via a direct network connection, thus ensuring a seamless change of operation.
Financial Services	Development/ Test & Manufacturing	A provider of financial software wants to utilize the benefits of the public cloud for his development activities and the manufacturing unit of the company. Nevertheless, he still wants to observe all compliance regulations, in order to keep control over the critical data of the application that is used to calculate the salaries.	To achieve this goal, 90% of all virtual machines (VM) were migrated from the internal IT infrastructure to Amazon EC2. This means that the major part of the development and production environment now runs within the Amazon public cloud. The data of the salary management application was transferred to a NetApp storage within a co-location data center. The development, as well as the production systems that run on Amazon EC2 machines, access the data over a direct network link. As a result, the financial services provider is able to meet the required compliance regulations despite the use of public cloud resources.
Research & Development	Cloud Bursting	A research institute needs to be able to handle peak loads that occur during high performance workloads, e.g. hurricane simulations, without having to invest in massive infrastructure resources beforehand. Total control over the data must be ensured at the same time.	The company still uses its own IT infrastructure. In addition, a dedicated fast MPLS link between the infrastructure and the co-location data center is installed. The simulation data or parts of it are kept in a NetApp Storage at the co-location. When the local computing resources have reached their limits, pre-configured Amazon EC2 instances with their relevant simulation applications are started up. The instances then access the data over an AWS direct connect link and carry out the calculations.

Other application scenarios are:

- Transcoding and provisioning of video content
- Enterprise File Sync & Share – similar to Dropbox
- Payment portal that meets compliance regulations
- Management of medical records
- Consolidation/ migration of data centers

5. Outlook

Digitization increasingly gains pace. Company executives and IT decision-makers are being called upon to become part of this development and to transform their companies or, more specifically, IT organizations. In this context, public cloud infrastructure offers ideal conditions through providing easy access to the IT resources needed to implement companies' business and IT strategies. The challenge is that due to the characteristics of shared infrastructure, compliance regulations as required by law often cannot be met. In addition, having authority of control over the company's data is of great importance. Data, after all, are regarded as a critical asset with a high business value in the digital age.

Being implemented via use of "NetApp Storage for Cloud", hybrid storage concepts offer an opportunity to resolve these challenges. Even today, IT architects no longer deal with a single cloud provider but rather seek diversity in multi-cloud environments to minimize the technical risks and benefit from cost advantages. To do so, they consequently need to ponder the questions of latency, in order to provide sufficient performance to their applications and workloads.

Hybrid cloud and multi-cloud architectures offer technical and organizational answers to these and other questions and place control over data back into the hands of the company. Therefore, in years to come, innovative enterprise storage concepts that fulfill the requirements of a hybrid and multi-cloud world will play an essential role in the pending transformation of many large and medium-sized companies.

About NetApp

Since our beginnings we have been pursuing the goal to put our customers on the road to success. In 1993 we delivered the first storage system to Patrick Mulrone, who back then worked as a system administrator at Tandem Computers. At that time, the so-called “Network Appliance” was employed to prevent the uncontrolled growth of the company’s infrastructure. He was thrilled by the simplicity of the solution.

Today – more than 20 years later – NetApp is still creating innovations that give our customers a competitive edge. From the boundaries of the universe (with the Large Hadron Collider at CERN) to the limits of what the human body can withstand (in Formula 1), we support our customers with solutions and speeds they would never have believed to be possible.

Leading companies all over the world rely on NetApp software, systems and services for managing and storing their data. We help companies and service providers with the planning, provisioning and further development of IT environments. Our customers also benefit from the cooperation with other leading technology providers. This allows for custom-made solutions that are specially tailored to specific needs.

The NetApp Team attaches great importance to the success of its customers. Our company culture and work environment are built around this premise. Together with our partners all over the world, we share the common goal of putting our customers on the road to success.



NetApp™

NetApp Deutschland GmbH
Sonnenallee 1
85551 Kirchheim bei München

TEL +49 30 887062201

WEB www.netapp.com

Author



Rene Buest

Senior Analyst & Cloud Practice Lead

rene.buest@crisp-research.com

Rene Buest is Senior Analyst and Cloud Practice Lead at Crisp Research, with focus on Cloud Computing, IT infrastructures, open source and Internet of Things. He is a member of the worldwide Gigaom research analyst network, a top Cloud Computing blogger in Germany and one of the world's top 50 bloggers in this area. On top of this, he is one of the world's top Cloud Computing influencers and one of the top 100 Cloud Computing experts on Twitter. For more than 16 years, he has been concentrating his work on the strategic use of information technology in companies. In addition, he analyzes the influence of information technology on our society and assesses disruptive technologies.

Rene Buest is the author of numerous professional articles on Cloud Computing and technology as well as being a lecturer and member of expert panels. He writes about cloud computing, IT infrastructures, technologies, management and strategies on CloudUser.de. He holds a master's degree in Information Science from the Bremen University of Applied Sciences, as well as an M.Sc. in IT Management and Information Systems from the Paderborn University of Applied Sciences.

About Crisp Research

Crisp Research AG is an independent IT research and consulting company. Backed up by a team of experienced analysts, consultants and software developers, Crisp Research analyzes current and future technology and market trends. Crisp Research supports companies with the digital transformation of their IT and business processes.

The assessments and comments by Crisp Research are published and discussed by numerous financial magazines, specialized journals for the IT business and social media. As contributing editors for leading IT publications (Computerwoche, CIO, Silicon et al.), BITKOM enthusiasts and sought-after keynote speakers, our analysts not only actively contribute to the debates about new technologies, standards and market trends, they are also among the relevant influencers of the industry.

Crisp Research was founded in 2013 by Steve Janata and Dr. Carlo Velten. The company focuses its research and consulting services on “Emerging Technologies” such as Cloud, Analytics or Digital Marketing and their strategic and operational implications for CIOs and decision-makers in companies



Weissenburgstrasse 10
D-34117 Kassel
TEL +49 561 2207 – 4080
FAX +49 561 2207 – 4081

E-MAIL info@crisp-research.com

WEB crisp-research.com
crisp-analytics.com

TWITTER [@crisp_research](https://twitter.com/crisp_research)

Copyright

Created on Behalf of:

NetApp Deutschland GmbH
Sonnenallee 1
85551 Kirchheim bei München

TEL +49 30 887062201
WEB www.netapp.com

All rights to this material are the property of Crisp Research. The data and information remain the property of Crisp Research AG. This work may not be reproduced, in whole or in part, without the prior written permission of Crisp Research AG.

Design, layout & infographics
Hellwig & Buntenschuch
E-MAIL info@hellundbunt.de
WEB hellundbunt.de

Weissenburgstrasse 10

D-34117 Kassel

TEL +49 561 2207 – 4080

FAX +49 561 2207 – 4081

E-MAIL info@crisp-research.com

WEB crisp-research.com crisp-analytics.com

TWITTER [@crisp_research](https://twitter.com/crisp_research)

