



# Cloud Data Fabric: Enterprise Storage Services in the Cloud

A Strategy Paper created on behalf of NetApp Deutschland GmbH

by Rene Buest, Senior Analyst & Cloud Practice Lead

## Executive Summary

---

- The maturing public cloud infrastructure gains increasing importance as an attractive alternative to on-premise enterprise IT infrastructure. Public cloud infrastructures offer an abundance of possibilities to CIOs when it comes to operating existing infrastructure and application environments more flexibly and at a lower cost.
- The existing public cloud storage services have been developed with a focus on the new generation of applications. This is why they are less well prepared to run existing enterprise applications.
- At present, the requisite storage concepts, standards and technologies for use of legacy enterprise applications on public cloud infrastructure are still not available.
- In order to ensure the continued existence as well as the proper operations of conventional application architectures on public cloud infrastructures, it is necessary to transfer the established and widespread standards into the public cloud.
- As legacy applications still constitute the lion's share of potential cloud migration candidates for new types of cloud native applications, well-known storage concepts are required to ensure the migration of existing applications without modifications to a public cloud infrastructure, where they continue to be operated.

## Table of contents

---

<b>1. Enterprise Storage: Fit for the Cloud?</b>	<b>4</b>
<b>2. Cloud Storage-Services: Enterprise-ready?</b>	<b>6</b>
<b>3. Cloud Data Fabric: Storage Management in the Cloud</b>	<b>8</b>
<b>4. Outlook</b>	<b>10</b>
<b>About NetApp</b>	<b>11</b>
<b>Author</b>	<b>12</b>
<b>About Crisp Research AG</b>	<b>13</b>

# 1. Enterprise Storage: Fit for the Cloud?

---

**Information.** A growing number of companies currently place the assessment of data in the center of their existing and new business models. The digital transformation leads to a new class of data-based products and services.

Apart from mobile devices and social media services, the services, sensors and devices from the Internet of Things (IoT) and the Industrial Internet furthermore generate huge data amounts, which are collected and analyzed over Big Data and Analytics solutions. For example, the Claas company, a leading manufacturer of agricultural machinery, and its partners KWS Saat AG and Lamken benefit from the IOT analytics solutions, as data services for farmers are offered on a common marketplace.

The basic challenge is this: in a controlled fashion, all corporate data must be stored on a medium that meets the requirements with respect to reliability, fail-safe operation and low fault tolerance, so that the fulfillment of possible compliance and governance regulations can be ensured.

Enterprise storage solutions that aim to meet these requirements have been developed over the past 10 years. These solutions are predominantly deployed for typical corporate scenarios, for example:

- Data Storage
  - Data Backup
  - Data Archiving
  - Disaster Recovery
- 

Although conventional enterprise storage solutions fulfill the requirements of these scenarios, they nevertheless disregard changes towards data-driven business models, as well as new services that are being created by IoT and the Industrial Internet. For this purpose a new storage class is required, one that behaves more flexibly and proves to be more scalable, satisfying the needs of both existing and new workloads.

With the advent of the cloud, more and more CIOs transfer their existing enterprise applications to the cloud, in order to benefit from its cost advantages and dynamic nature. At the same time, the associated data must also be relocated accordingly, leading to hybrid architectures that pose a challenge for the technology and data management. On the one hand, data have specific inertia<sup>1</sup> that makes the cloud transfer difficult.

On the other hand, widely used conventional enterprise NAS/-SAN-storage concepts cannot be implemented in the cloud by means of technology standards like iSCSI or protocols like NFS.

---

<sup>1</sup> The cloud - an Integral Part of All Business Processes, Crisp Research AG, April 2014, <http://www.crisp-research.com/report/die-cloud-fur-alle-falle-datenkontrolle-und-globale-skalierbarkeit-mit-hybrid-und-multi-cloud-architekturen-sicherstellen/>

The reason: public cloud providers prefer object and block storage concepts, in order to store data in highly redundant, distributed and scalable ways. Enterprise applications, which previously were used on on-premise infrastructures and are now being migrated into the cloud, cannot handle this type of data storage.

Therefore, a few fundamental questions need to be answered prior to the cloud migration of legacy enterprise applications and their corresponding data assets:

- Are there enterprise functions in the cloud to safeguard integration with the on-premise storage management?

---

- Is it possible to easily move the data and can they be used by migrated company applications?

---

- Does the cloud provide the necessary functions and protocols, for example iSCSI, NFS, CIFS, to implement existing storage scenarios?

---

- Can existing tools and interfaces still be used?

---

## 2. Cloud Storage-Services: Enterprise-ready?

---

In order to get an understanding as to whether the cloud migration of existing enterprise applications and their corresponding data can be carried out easily, some background knowledge of public cloud storage services is necessary.

Public cloud providers generally rely on three types of storage concepts:

→ **Instance Storage**

Instance storage is a temporary local mass storage system on a virtual machine (VM), where VM-centric applications and data can be placed. The heart of this storage is local memory. This means that as soon as the virtual machine is shut down or fails, all data will be lost.

---

→ **Object Storage**

Object storage organizes the data without hierarchical levels, within a flat address space in a storage pool. As a means of identification each object is assigned a unique feature that is stored in its metadata. Therefore, servers, services and applications don't need to know the physical storage location of the data - instead they simply access the object via its metadata.

---

→ **Cloud Block Storage**

Block storage is organized in volumes (blocks). Here each block represents an independent hard drive that is managed by the operating system. Block storage can be assigned to a virtual machine, to store data there on a long-term basis.

---

Public cloud providers have a selective choice of storage types that can be tailored to fit application-specific needs. Yet, cloud-focused concepts of this kind have downsides when it comes to usage in the context of legacy enterprise applications. These include:

- Instance Storage is only suitable for temporary data storage.

---

- Access to object storage is enabled through REST APIs, with each access being recorded in the program logic of newly developed cloud-based applications.

---

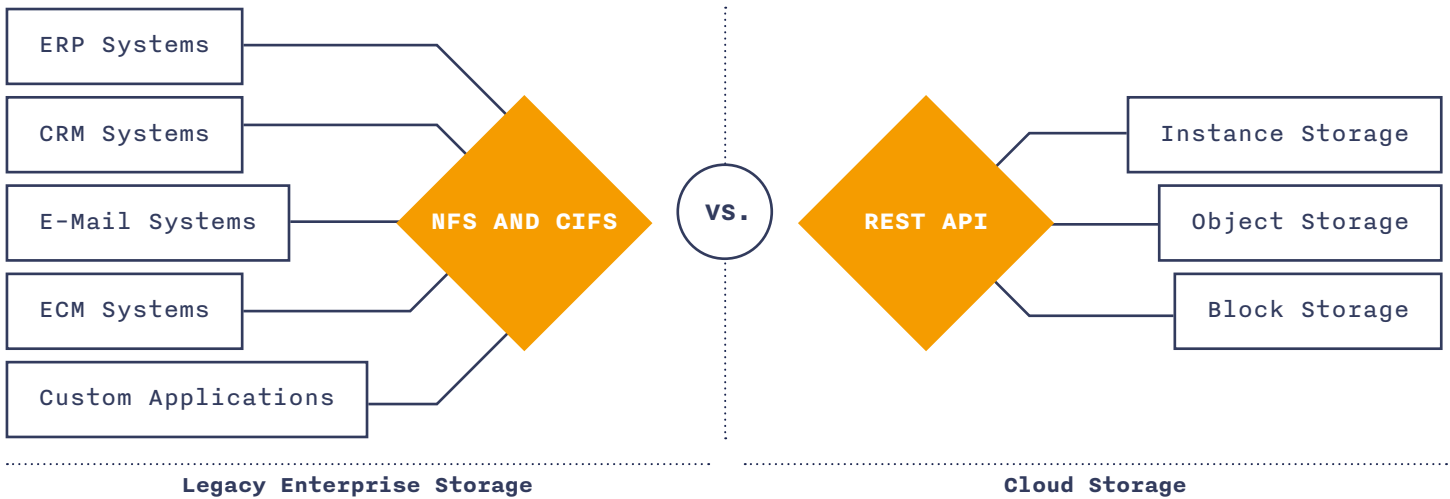
- Non-cloud ready application architectures of enterprise applications in most cases need NFS/CIFS-storage, as they do not work with object storage.

---

- Block storage can only be utilized in conjunction with a virtual machine (VM). Central access by multiple VMs is limited.

---

Legacy Enterprise Storage vs. Cloud Storage



QUELLE:  
Crisp Research AG, 2015

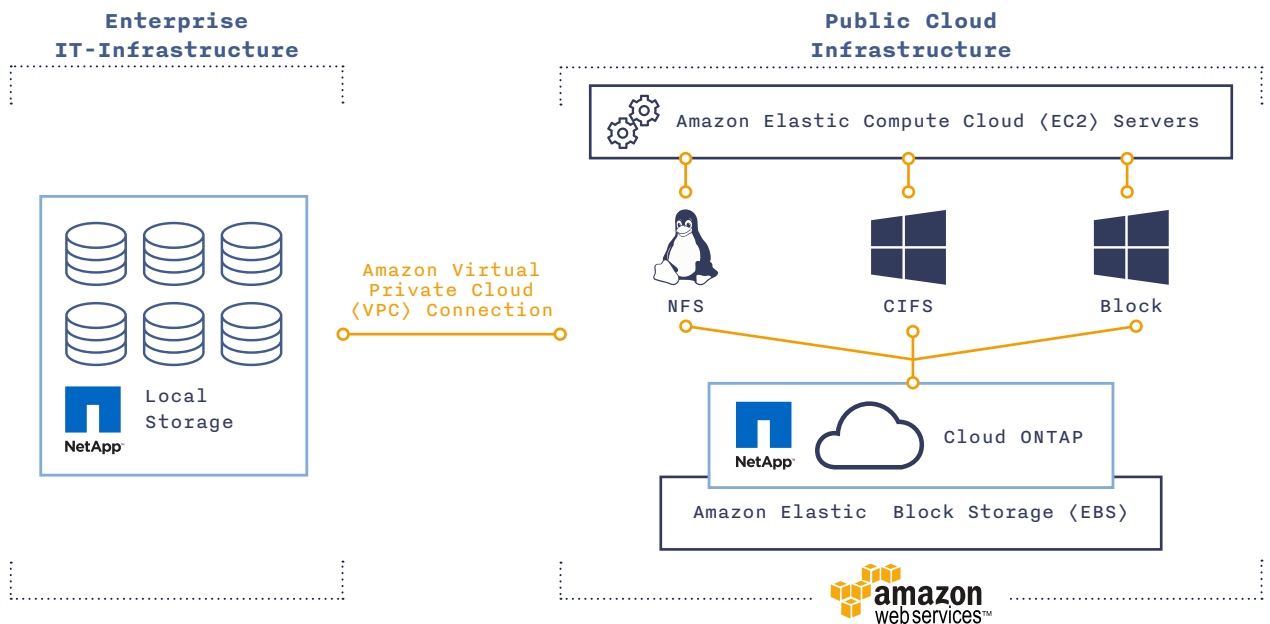
Current public cloud storage concepts have been developed primarily with the focus on novel or rather cloud native applications. They are actually not meant to be used together with existing enterprise applications. This becomes particularly evident when it comes to compatibility with established standards and data mobility<sup>2</sup> (data gravity) during migration, as well as to the subsequent level of control.

<sup>2</sup> Ensuring Data Control and Global Scalability with Hybrid Cloud and Multi-Cloud Architectures, Crisp Research AG, 2015, <http://www.crisp-research.com/report/die-cloud-fur-alle-falle-datenkontrolle-und-globale-skalierbarkeit-mit-hybrid-und-multi-cloud-architekturen-sicherstellen/>

### 3. Cloud Data Fabric: Storage Management in the Cloud

Public cloud infrastructure gradually turns out to be a suitable alternative to on-premise enterprise IT infrastructure. As a result, companies have various possibilities for making their existing IT environment more flexible and more cost-efficient. Nevertheless, public cloud infrastructure lacks the requisite storage concepts, standards and technologies for the deployment of legacy enterprise applications. In order to safeguard the survival of existing applications and to ensure that the necessary requirements are met, NetApp offers the “Cloud ONTAP” concept. For the implementation of this enterprise cloud storage scenario, NetApp is presently cooperating with the public cloud provider Amazon Web Services.

**Enterprise Cloud Storage Scenario**  
**NetApp Cloud ONTAP on Amazon AWS**



QUELLE:  
 Crisp Research AG, 2015

“Cloud ONTAP” represents a special storage operating system. Provided in the form of a „Virtual Appliance“, this system is installed on an “Amazon EC2” instance (virtual machine) as a software. Later, Cloud ONTAP acts as a storage controller, taking care of the data management within a virtual environment on the Amazon Cloud. This virtual environment, which is built up and operated by a company, constitutes the basis for the implementation of public cloud NAS and SAN Storage scenarios for legacy enterprise application environments and supports industry standards like NFS, CIFS and iSCSI. All data are managed by Cloud ONTAP and placed on Amazon Elastic Block Storage (EBS) volumes in an encoded form. The “onCommand Cloud Manager“ supports administrators with their management tasks, provides an overview of the current storage consumption by the user and logs how many times data access has taken place. For this reason, companies have the required control mechanisms at their disposal to ensure the fulfilment of governance and compliance regulations.



If a “Clustered Data ONTAP” is deployed within a company’s own IT infrastructure, hybrid scenarios can also be implemented, allowing for the further expansion of enterprise storage services in the cloud. This way, for example, the storage of a legacy enterprise application can be separated and outsourced to a more cost-effective public cloud storage. The application can access the data in the public cloud from within the on-premise environment via a VPN link (Amazon Virtual Private Cloud (VPC)) or a dedicated “Direct Connect” connection.

## 4. Outlook

---

**As public cloud infrastructure matures, it is being increasingly singled out by CIOs and infrastructure decision makers when it comes to migration planning for existing enterprise applications.** Here all parties involved face the drawback that conventional legacy applications are not designed to be operated on cloud infrastructure in the first place. To make things worse, these non-cloud ready application architectures rely on protocols and technical standards that are currently not supported by the public cloud storage services of the providers.

New storage concepts are required during and after migration, to ensure the continued existence of conventional enterprise applications on cloud infrastructure while supporting the availability of established and widespread standards in the public cloud. Since many applications inherently do not support object storage, NAS solutions are among the favored options to run legacy enterprise applications in the public cloud.

Taking into consideration that a large number of legacy enterprise applications will be migrated into the cloud in the future, storage concepts like NetApp Cloud ONTAP are the solutions of choice for operating existing applications on public cloud infrastructure, while preserving their complete functionality and structure.

## 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 89 900 594 0

FAX +49 89 900 594 99

WEB [www.netapp.com](http://www.netapp.com)

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

## Author

---



### Rene Buest

Senior Analyst & Cloud Practice Lead

[rene.buest@crisp-research.com](mailto: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.** Prior to this, he was principal analyst at New Age Disruption and a member of the worldwide Gigaom Research analyst network. René Büst is 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 and Google+. As of the mid-90s, 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. He also regularly writes for renowned publications like Computerwoche, CIO Magazin, LANline and Silicon.de and his contributions are cited by German and international media – among them The New York Times, Forbes Magazine, Handelsblatt, Frankfurter Allgemeine Zeitung, Wirtschaftswoche, Computerwoche, CIO, Manager Magazin and Harvard Business Manager. In addition, René Büst is a lecturer and member of expert panels. He is the founder of CloudUser.de and writes about cloud computing, IT infrastructures, technologies, management and strategies. René Büst 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 AG

---

**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](mailto:info@crisp-research.com)

WEB [crisp-research.com](http://crisp-research.com)  
[crisp-analytics.com](http://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 89 900 594 0

WEB [www.netapp.com](http://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 & Bunttenbruch

E-MAIL [info@hellundbunt.de](mailto:info@hellundbunt.de)

WEB [hellundbunt.de](http://hellundbunt.de)

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

**E-MAIL** [info@crisp-research.com](mailto:info@crisp-research.com)

**WEB** [crisp-research.com](http://crisp-research.com) [crisp-analytics.com](http://crisp-analytics.com)

**TWITTER** [twitter.com/crisp\\_research](https://twitter.com/crisp_research)

