Move Applications not Data – A new Paradigm for the Future Internet

Peter Domschitz, Markus Bauer, Jürgen Sienel and Marcus Kessler
Bell Laboratories, Service Infrastructure Research
Alcatel-Lucent, Stuttgart, Germany
{Peter.Domschitz, Markus.Bauer, Juergen.Sienel, Marcus.Kessler}@alcatel-lucent.com

I. INTRODUCTION

The Internet is changing the way how we consume media. Enabled by Internet technologies, we are quickly moving towards a point allowing everybody to enjoy media services like 3-D broadcasting, live as well as time-shifted, whenever wanted, wherever needed on any preferred device. Even more the user will not only be a spectator but an immersed participant. The Web is the catalyst for a whole class of new personalized media centric applications, like massive sharing of multimedia content even in real-time or immersive multimedia communication. These services will not be realized as a pure content stream, but as an orchestrated flow of media processing functions that will provide the requested data at the appropriate time, location and format. With the introduction of HD video formats, the transferred data volume will outrun the size of the code building in the data transformation process. Therefore intelligently placing service components on a distributed service infrastructure will provide a way to increase the scaling of the Internet infrastructure in the future. We will motivate the concept of "Move Apps not Data" and elaborate some enablers requested from future Internet research.

II. THE "MOVE APPS NOT DATA" PARADIGM

Up to now increased transport capacity demands in the networks were mainly achieved by enhancing installed bandwidth either by technological break-troughs or installation of new infrastructure elements. But there are substantial concerns, that this cannot be expected to last, at least at reasonable costs [1][2]. As future network enhancements become more and more challenging, there is need for alternative ideas. A common approach to achieve this goal is to add "higher layer" intelligence to the networks, aiming to reduce overall traffic, thus enhancing available transport capacity by localizing traffic. The first huge success of this concept was the invention and successful introduction of Content Delivery Networks. CDNs basically enabled the massive scale adoption of (media) services comprising broadcast delivery characteristics in the Internet. But today we see an emerging trend towards personalized media streams, where such streams need to undergo processing somewhere in the Internet, i.e. in the cloud, enabling the evolution of e.g. IP-TV towards personalized "multi-view" video services [3] or cloud based gaming services like "onlive" [4]. As CDNs are build for the efficient delivery of the same content to a multitude of receivers, this trend towards individual content streams requiring processing in the network is going to challenge the future Internet infrastructure.

Today's applications and corresponding cloud infrastructures are typically designed in a way that data is moved through the network to a dedicated location (i.e. data center) where the application is executed. Preserving that cloud computing paradigm would result in huge amounts of traffic which needs to be transported to "arbitrary" data centers where the processing is located. We believe that this paradigm will change in the future, meaning that an intelligent infrastructure will also force the movement of applications or parts of it next to the data. Such "Move Apps not Data" schemes can offload unnecessary "long distance" traffic from the networks by localizing traffic and thus will help to overcome the issue of limited availability of transport capacity in future networks.

The following simple example should illustrate the basic idea behind the "Move Apps not Data" (MAD) scheme: We assume a video rendering service like Animoto, which can produce orchestrated video clips from your personal pictures and music. A user provides the video rendering engine with a reference to the original content, stored at its preferred storage provider network (i.e. data center) somewhere in the cloud. In a next step the rendering engine will download the referenced files and render the video clip from this data. As a consequence huge amounts of imaging and music data needs to be transferred between the, most likely, different locations where the data is stored and the video rendering is executed. As stated above The "MAD" concept aims to avoid such unsustainable traffic patterns in future networks by localizing traffic. In case of our video rendering example the overall traffic imposed on the network can be dramatically reduced by moving and executing the video rendering on cloud processing resources, which are located very next to the user's data.

This motivates us to elaborate new ways to optimize the delivery of (real-time) media services on top of distributed cloud environments as the key scenario for "Move Apps not Data".

III. SERVICES IN THE FUTURE INTERNET

Even with today's technology, offloading of computing infrastructure into the Internet has become a commodity. Cloud providers, like Amazon EC2 or Microsoft Azure offer their infrastructure or platforms as a service, providing features like automated scalability and instant deployment, which supports the very dynamic needs of Internet services like Facebook or Animoto. But this burdens significant cost on the networks:

The overall core network load is increased because instead of keeping traffic local more traffic is routed to centralized data centers that process the data and send it back to the requestors.

While this seems still feasible for traditional request/response based web-services, it might break the design of today's Internet architecture for massive media centric applications.

In future the Internet will embed intrinsic capabilities to directly support a service oriented computing paradigm that enables developers and end-user to execute their personalized applications onto an integrated network and computing infrastructure. Applications are then build from loosely coupled service components which are no longer bound to a specific host but become moveable objects, which can be dynamically deployed on the distributed computing resources and collocated to the data-flows between the sources and sinks. Although from the user's perspective the cloud behaves like a centralized server, it utilizes an aggregated or distributed set of free resources in a coherent manner.

Applications with varying service characteristics can profit to a different extend from the MAD concept. Huge benefits can be achieved on applications that require a consistent flow of continuous data over a certain period or require the transfer of huge amounts of data for processing. But one can easily image other examples, which require only very limited transfer of data. In such cases service transfer overhead and instantiation cost might exceed the gained benefits. As a consequence within the MAD concept there is need to define mechanisms allowing the retrieval of "meta-information" associated with data, e.g. where is the data located, how much data needs to be transferred for service execution, is the data a constant media (e.g. video) stream or just a bunch of data which must be transferred prior to service execution.

IV. NETWORKING ASPECTS

To support media cloud scenarios inherently from the network architecture, some basic principles from the existing Internet architecture need to be reconsidered. First, we need to extend well-known principles from content networking to support our "Move Apps not Data" approach. Content networks explore locality of data, i.e. instead of serving a request for a data at the source, a local cached copy is delivered. Van Jacobsen has proposed a scheme that directly addresses the content and uses this information for routing, instead of IP addresses representing the location where it was generated [6]. A natural extension in the MAD context would mean, that we do not only address the content, but also can address a service that is able to provide the requested data and instantiate a processing pipeline to do the necessary transformations. Instead of centralized processing for all users in a single domain, media flows are combined or split at appropriate location exploiting intrinsic "multi-cast" capabilities in the network layer where available. This is beneficial against existing schemes, where multicast has to be explicitly incorporated by the service developer not knowing if it is supported in the network and thus can only be achieved by means of overlay mechanisms.

In section III we introduced the notion, that MAD requires knowledge on "meta-information" to ensure educated service placement in the cloud. If we are able to accurately predict from this meta-information the traffic patterns of (media) flows exchanged between different service components, a future MAD enabled network can operate on such flow patterns directly, instead of executing routing decisions on a per packet basis. Thus MAD can enable efficient flow based switching by providing available meta-information of media streams, exchanged between MAD service components to the network. This information will enable the control plane in future networks to increase the overall throughput. But such schemes also need to consider, that a flow based switching paradigm is achieved at the cost of supporting more dynamicity in flow control handlers. To limit such costs future networks need to provide capabilities to aggregate multiple of such streams. By introducing this aggregated granularity of joint streams between the distributed (micro) data centers the control complexity in the core network itself can be limited. An equally important requirement is, that flows are now handled in the network in such ways, that uninterrupted relocation of media flow end-points, which are no longer machines but services (i.e. MAD service components) is supported. In consequence client APIs like the socket interface have to be rethought. As MAD enabled services are built from self-contained components generally operating on input stream(s) of data to generate their output data, which is then distributed to the subsequent consumer components of this service, the use of dedicated sockets for communication purposes is no longer sufficient and new paradigms need to be considered in the context of future Internet.

V. CONCLUSION AND OUTLOOK

"Move Apps not Data" proposes a different approach to the way services might be realized in the future Internet. By taking advantage of a distributed cloud infrastructure and new networking concepts we enable service developers to efficiently realize distributed applications without the burden of handling service placement and connectivity inside their software. Although this requires new service infrastructure and middleware technologies there is significant value in localizing network traffic by breaking the traditional client-server model. Although realization through overlays is a valid starting point, we consider an integrative approach with future Internet networking technologies as beneficial. A simplified API towards the network layer eases the building of a global infrastructure.

Alcatel-Lucent will prototype and evaluate the potential benefits of MAD in the context of G-Lab project NETCOMP [7].

[1] A.D. Ellis, J. Zhao, D. Cotter: "Approaching the Non-Linear Shannon Limit," Journal of Lightwave Technology, vol. 28, no. 4, 2010.

[2] S. Namiki, T. Hasama, H. Ishikawa: "Challenges for the future networks and enabling photonic technologies," OptoElectronics And Communications Conf., Hong Kong, July 2009.

[3] I. Ahmad: "Multi-View Video: Get Ready for Next-Generation Television," IEEE Distributed Systems Online, vol. 8, no 3, 2007..

[4] OnLive, http://www.onlive.com/

[5] J. Domingue, D. Fensel: "Toward a Service Web: Integrating the Semantic Web and Service Orientation", IEEE Intelligent Systems, 23, Jan./Febr. 2008, pp. 86-88.

[6] V. Jacobson, et.al.: Networking named content. Proceedings of the 5th ACM International Conf. on Emerging Networking Experiments and Technologies (CoNEXT 2009), Rome, Italy, Dec. 2009.

[7] BMBF Förderkennzeichen 01BK0940, www.germanlab.de/ .