09072 Abstracts Collection
Bandwidth on Demand
— Dagstuhl Seminar —

Panayotis Antoniadis[1], David Hausheer[2], Kohei Shiomoto[3], Burkhard Stiller[2] and Jean Walrand[4] (Eds.)

[1] Université Pierre et Marie Curie, Paris, France
Panayotis.Antoniadis@lip6.fr
[2] University of Zurich, Switzerland
[hausheer,stiller]@ifi.uzh.ch
[3] NTT Network Service Systems Labs, Tokyo, Japan
shiomoto.kohei@lab.ntt.co.jp
[4] University of California, Berkeley, USA
wlr@eecs.berkeley.edu

Abstract. From February 8-11, 2009, the Dagstuhl Seminar "Bandwidth on Demand" was held in Schloss Dagstuhl – Leibniz Center for Informatics. During this seminar, participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the talks given during the seminar are put together in this paper.

Keywords. Bandwidth on demand, bandwidth provisioning, bandwidth trading, network economics, resource allocation, network virtualization, wireless mesh networks, peer-to-peer networks, business modeling, telecommunications, market mechanisms, regulation, legislation

09072 Executive Summary – Bandwidth on Demand

David Hausheer (University of Zurich, CH)

This paper is the executive summary of the Dagstuhl Seminar 09072 "Bandwidth on Demand" which was held from February 8-11, 2009, in Schloss Dagstuhl – Leibniz Center for Informatics.

Keywords: Bandwidth on demand, bandwidth provisioning, bandwidth trading, network economics, resource allocation, network virtualization, wireless mesh networks, peer-to-peer networks, business modeling, telecommunications, market mechanisms, regulation, legislation

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/2132

Bandwidth Externalities and QoS Growth: A Long-Run Economic Approach

Athanassios Androutsos (Athens University of Economics and Business, GR)

In this talk, we investigate the role of bandwidth as Quality-of-Service (QoS) growth engine in the Internet economy. Bandwidth has implications for the service quality because it is used as an input in the QoS production process. The main results of our approach are that bandwidth exhibits substantial positive externalities since decisions on bandwidth levels made by a single user, affect all other users' output, that users cannot internalize bandwidth externalities in competition thus leading QoS to lower rates of growth than command economy and that the increase of bandwidth per capita will increase the rate of QoS growth.

Keywords: Bandwidth growth, QoS growth, Internet growth, network economics, network externalities

Resource Allocation and Provision in Non-profit Network Facilities

Panayotis Antoniadis (Université Pierre et Marie Curie, Paris, FR)

Distributed network facilities, like Planetlab, rely on the individual contributions of participants in order to build an infrastructure of sufficient size to allow them to run large scale experiments on overlay networks created over the Internet. In this presentation, I discuss our on-going work on the design of resource provision and allocation policies in this context, arguing in favor of a rule-based mechanism relating contribution with allocation as a realistic and effective way to increase the system's efficiency.

Similar policies are required for the case of federation agreements between network facilities run by independent authorities. Such a federation is already in progress between Planetlab Central and Planetlab Europe which currently employs a simple "peering agreement" as its resource management policy. But the problem will become more challenging as more diverse facilities (*e.g.*, Federica, wireless testbeds, etc.) are getting involved and demand is rising. I present our first thoughts on how to design more sophisticated federation policies that will provide the appropriate incentives to independent facilities to agree to federate and share their resources efficiently.

Keywords: Federation policies, network facilities

Analysis of Aggregation Strategies in Bandwidth-on-Demand Markets

Fernando Beltran (University of Auckland, NZ)

Some features of the current/future telecommunications landscape are causing a shift in the way consumers and providers access other providers' networks; in particular, uncertainty, brought in by the changing consumer communications needs, makes network capacity provisioning more volatile. If bandwidth is to be purchased on an "on-demand" basis then parties interested on having access to communication networks might have to abide by new types of transactions – at least new for the industry – which would certainly include the use of auctions and more complex bidding processes. Assuming that very important technical aspects regarding the feasibility of network access points are in place, in our work we are particularly interested on how a buyer, that is, individual consumer, network reseller or aggregation agent, would strategically decide to aggregate bandwidth purchased from more than one bandwidth provider. It's hard to envision which part of the network will see such arrangements developing, either wholesale or retail, but they most likely will appear as a response to bandwidth on demand markets.

In summary, this work studies efficiency and equilibrium properties of proposed bidding strategies that aggregate purchase of network resources such as bandwidth. Such approach becomes relevant to bandwith-on-demand markets as users may respond to changes in the market environment by strategically expressing their their network resource preferences, while providers can more agily decide who gets what and the price to be paid.

Keywords: Auctions, multiaccess markets, aggregating bidding strategies, competitive on-demand markets

QoS Support for Overlay Multicast

Torsten Braun (Universität Bern, CH)

Most applications used today such as Voice over IP (VoIP), watching movies, and even multimedia streaming are not very bandwidth demanding or do not require hard Quality-of-Service (QoS) guarantees. Moreover, the way we consume content (*e.g.*, podcasts) and communicate (chat, email, VoIP) nowadays allows rather using mechanisms like caching, buffering, multicasting instead of resource reservation mechanisms to achieve reasonable application level quality. An overlay multicast solution has been developed that can operate on top of resource reservation as well as measurements based QoS mechanisms in order to support heterogeneous QoS scenarios. The overlay multicast solution supports services such as real-time streaming, games, and Internet television (IPTV).

Keywords: QoS, multicast, overlay networks

Multihoming in Heterogeneous Wireless Access Networks

Georg Carle (TU München, DE)

Upcoming wireless access networks offer the capability of allowing to switch between heterogeneous wireless access technologies. Example access technologies are UMTS, LTE (Long Term Evolution) as being standardized within 3GPP, and wireless LAN as being standardized within IEEE 802.11.

One approach of switching between different access technologies is that of end systems equipped with multiple network interfaces that autonomously decide which wireless access technology to use. Another approach is that of a wireless service provider with network resource management entities that allow to decide in which way terminals, or traffic flows, can be switched from one wireless access network to another, thereby possibly switching access technologies.

These scenarios represent several challenging issues, such as how to provide the network resource management entities with the required measurement information. In this talk, the approach of the Generic Measurement Infrastructure (GMI) is presented which uses a publish/subscribe system for collecting relevant information, and for offering this information to the entities that need them for resource management decisions. Open issues include identifying convincing solutions towards the instrumentation infrastructure, addressing (a) how flexible the instrumentation infrastructure should be, (b) which are the requirements of entities in need of measurement information, and (c) whether and how to restrict access to the measurement information. Additional open issues are related to the decision process, in particular about (a) whether intelligence for handover should be rather implemented into the end node, or into the network, (b) which is the impact of tariffs onto decisions, and (c) how to optimize tariffs.

Keywords: Multihoming, wireless access, handover

Value of Reputation Management for BoD?

Zoran Despotovic (DoCoMo Euro-Labs – München, DE)

In this talk I investigate the potential value of (possibly decentralized) reputation management to bandwidth on demand applications. The setting can be that of many users on a shared medium such as a wireless cell, where one user can give up its medium usage in order to provide better conditions for another user. Frequent dealing with strangers and low stakes of transactions may require soft assurance mechanisms such as reputation. I wonder if such settings are of interest for BoD and, if yes, what reputation mechanisms are best suitable.

Techno-Legal Bandwidth on Demand Perspectives

Jochen Dinger (Universität Karlsruhe, DE)

The topic of "Bandwidth on Demand (BoD)" might lead to new innovative systems in the future. However, it is questionable whether long established dominant market players might discriminate such systems because of their existing business models. Thus, we analyze from a legal perspective if there are market regulations that could ensure the protection of new innovative systems. Our analysis is based on peer-to-peer (P2P) systems as for a legal analysis concrete facts about a system are necessary and P2P systems are similar and/or a basis for BoD respectively. First, we show the innovative potential of two new P2P systems. Afterwards, we provide a classification of the most relevant European directives, laws, and legal opinions on electronic communication systems. Based on that we develop a new classification scheme and show that virtual networks like P2P networks are subject to be networks in the legal sense. Hence, telecommunication market regulations can also be applied to protect P2P systems of being discriminated. Finally, we state some open questions with respect to the broader context of network neutrality.

Keywords: P2P, telecommunication market regulation, network neutrality

Convergence without Conflation

Adrian Farrel (Old Dog Consulting – Llangollen, GB)

Implicit in the discussion of bandwidth on demand is the concept of network convergence. Client networks are brought closer together as they share the same physical resources within the server network, and client and server networks are more tightly integrated to give the client the ability to dynamically request capacity across the server. Changes in the commercial world are forcing network operators to consider new revenue streams, and network virtualisation is rapidly rising on the list of potential ways to grow income from existing transport resources. At the same time, radical changes in network usage (from IPTV to data-capable mobile devices) are forcing the operators of customer networks to look for short-term, high-capacity connectivity from their transport providers. Both of these ideas drive the requirement for dynamic core transport networks that can offer bandwidth on demand.

But these concepts of convergence do not sit comfortably with every service provider. For some, the viability of bandwidth on demand is not consistent with their network topology. For them, dynamic capacity could only be provided by dedicated physical topology and so they describe the concept as "truck roll on demand". Clearly, there is a need for these providers to re-plan their networks if they are to benefit from network convergence. For other service providers there is tremendous unease about allowing traffic demands in a client network to cause

the allocation of transport network resources. The fear is that high-capacity core connections will be set up to support micro-flows, and that transport resource usage will flap in response to short-lived user demands. In effect, service providers are afraid that network convergence will result in conflation of network operation.

To combat these issues we have invented two significant functional components: the Path Computation Element (PCE), and the Virtual Network Topology Manager (VNTM). I propose to show how these components can be used to support bandwidth on demand with full operator control of the way that the server network operates while still allowing dynamic, client-driven requests for bandwidth. I will also show how these components can be integrated with policy components, revenue and administrative function, and with the Resource and Admission Control Function (RACF) developed by the ITU.

Keywords: Path computation element, virtual network topology manager

Bandwidth on Demand – An ISP Point of View

Isabelle Hamchaoui (France Télécom R&D – Lanion, FR)

Bandwidth on Demand is not a new concept. It has been studied in detail in the framework of ATM networks, for both residential and business markets. Lessons from the past show that a real need for bandwidth on demand is highly questionable. Indeed, business and residential customers are much more searching for Quality of Experience (QoE), and bandwidth is only one factor impacting quality. Allocating bandwidth on a statistical manner, associated with statistical QoS guarantees is probably enough for the vast majority of customers. However, some specific needs for BoD based on resource reservation and deterministic QoS may exist in some minority cases, as research communities for example.

Keywords: Reservation, QoS, QoE

Two Challenges: "Wireless" and "As a Service"

Hannes Hartenstein (Universität Karlsruhe, DE)

Following J. Dinger's talk on our research in the field of P2P and legal issues, in this short presentation I looked at two challenges for bandwidth on demand. The first one, based on our research in wireless vehicular communication, addresses the issue of dynamically controlling the use of the scarce channel bandwidth for the benefit of the vehicular traffic system. This issue can be seen as a "non-standard" BoD challenge. The second issue raised in this talk is based on my experience as a director of SCC: while the BoD issues might be solved "asymptotically", they are "real" right now in the field of server farms, grids, and cloud computing. For example, flexible and quick reconfiguration of virtual networks is a big issue that is somewhere between research and business strategies.

Virtual Networks on Demand

David Hausheer (University of Zurich, CH)

Network virtualization allows to isolate different network applications from each other in a flexible manner. But more importantly, it allows to share physical network equipment such as fibers and routers in a transparent and cost-efficient manner. Especially, virtualization at the physical (optical) layer seems to be much more efficient in terms of energy consumption compared to routing at the higher layers (*e.g.* 25% of energy usage in the SINET3 research network in Japan), an important issue considering that energy costs are currently increasing.

The key question is what does "on demand" really mean? Does it mean reservation with certain guarantees for QoS? It is clear that customer requirements may change over time, especially for ISPs or large ASPs. For these types of customers it is necessary to enable dynamic allocation of bandwidth based on current application needs. An important question to be addressed here is the appropriate business model to be applied, in order to allow customers and providers to buy and (re-)sell bandwidth services on demand. Another important question is the time-scale, i.e. what are reasonable start intervals and durations of such services and whether it needs to be possible to reserve bandwidth in advance.

Keywords: Network virtualization, networks on demand

Considering Transaction Cost as a Performance Metric for Charging and Billing Systems

Christian Hoene (Universität Tübingen, DE)

This position statement proposes the use of transaction cost as performance metric for charging and billing of communication services. The concept of transaction cost has been introduced by Nobel Laureate and economist R.H. Coase. In communication networks, transaction costs include all expenses to inform the users, to search for providers, to negotiate the conditions of contracts between providers and users, to transfer money, to monitor the service provision, and to finalize the transaction. We present the design of a charging and billing system for bandwidth on demand that has very low transaction costs and we modeled its design tradeoffs analytically.

Keywords: Transaction costs, Internet charging

Energy On Demand – Lessons Learned from the Telecoms

George Huitema (TNO – Groningen, NL)

The Energy Business is changing dramatically. There are three main developments to be mentioned. First, energy is no commodity anymore since the changing climate, dependence on political unstable countries makes people aware that we have to use energy more efficient and effective by using more renewables, reduce CO2 emission and reduce the use of fossil fuels.

Secondly, the traditional hierarchical one-to-many distribution model of energy provision transforms to a model of many-to-many (also called P2P). This makes the balancing of demand and supply of energy key. Energy forecasting is very important since energy cannot be easily stored, which means that if there is a momentary shortage of energy supply immediately so-called virtual power plants – formed by local energy production, cogeneration – should fill in these gaps. Here smart meters will provide the actual usage details while market mechanisms will ensure that production meets the consumption.

Finally, on top of a physical energy infrastructure an adequate ICT infrastructure should enable the above mentioned transformation to a mixed central/decentral distribution network with a smart balancing of demand and supply (smart grid). Here, the energy world could learn from the Telecoms world: (a) the design of a suitable information architecture (demand, supply, usage); (b) pricing, incentive and billing experiences and functionalities; (c) demand-supply mechanisms (BoD).

Conclusion: the Energy world powers the ICT world but vice versa the ICT world has to "power" the Energy world.

Keywords: Energy, smart meter, billing

Bandwidth: No Demand

Aiko Pras (University of Twente, NL)

With the introduction of new technologies, such as lambda switching, it is technically possible to "sell" Bandwidth on Demand. Research networks, such as those collaborating within the context of the "Global Lambda Interconnection Facility" (GLIF), already offer dedicated Lambdas. Customers are, for example, CERN and VLBI researchers. Offered capacity varies from hundreds of Mbps, to several Gbps (in fact, the offered service is SDH/Sonet switching, and not real lambda switching).

The question is now whether normal end users would also be interested in buying bandwidth on demand. Such demand may emerge if the resource is scarce and has a clear value for the end user. Unlike oil, gas, food etc, bandwidth need not be scarce. In 10 years from now we can have an all optical Internet core, offering Pbps. Bandwidth may therefore be ubiquitous, instead of scarce. Second,

end users do not care about bandwidth, but they care about services (which, in turn, may use bandwidth). They will be willing to pay for these services, but not for the bandwidth in separation. If the service fails, but bandwidth is available, users do not want the bandwidth.

Conclusions are: (1) Bandwidth is cheap; (2) there will be enough for every one Bandwidth, once available, will be used; (3) but users don't care about bandwidth; (4) users care about applications; (5) users might hardly be willing to pay for bandwidth.

Keywords: Bandwidth no demand

The User Knows Best: A QoX-based View on *-o-D

Peter Reichl (FTW Vienna, AT)

BoD mechanisms usually focus either on an ISP-to-ISP or an ISP-to-user scenario. However, users are not interested in bandwidth per se (cf. the contribution by A. Pras), rather they would like to go for quality, which, for a long time, has been captured, *e.g.*, by the notion of "Quality-of-Service". However, whereas QoS in its early days was clearly directed to user satisfaction, over time this has evolved into a strictly techno-centric approach, investigating parameters rather than effects. Therefore, more recently the notion of "Quality-of-Experience" has appeared which we argue to provide an umbrella concept for BoD as well as other "*"-on-demand notions (cf. the contribution by B. Stiller). The resulting new perspective is illustrated in terms of a feedback-based charging scheme based on simultaneously signaling lack of user satisfaction and willingness-to-pay for better quality, thus realizing what could be called a "Quality-on-Demand" mechanism.

Privacy on Demand: A Difficult Challenge in a Highly Connected World

Giancarlo Ruffo (University of Torino, IT)

Nowadays many technologies are available for providing bandwidth to the consumer upon demand. At the same time, bandwidth via the broadband Internet is given even when the end user did not demand for it. Most of the many popular existing services are implemented using indiscriminate access to the Internet, even when alternative solutions can be addressed.

One motivation to look for alternative approaches to BoD is privacy, which deserves more attention during services design. During this talk, two case studies are provided: 1. the localization server used by the navigator installed in the iPod Touch (as well as in the iPhone), because no GPS is provided with the device; 2. recommender systems that compare the profiles and the preferences of all

the users in order to trigger personalized suggestions. There is evidence that epidemic algorithms, federated architectures, P2P systems, and so on, may be used to solve locally some of the related issues. Do centralized solutions always worth the sacrifice of privacy? Can we demand for bandwidth other than the one provided by world wide backbones?

Keywords: Social and legal impact, P2P, local access vs. internet access

AoD, BoD, CoD, ..., RoD, SoD, ..., and XoD

Burkhard Stiller (University of Zurich, CH)

Does the research community address in the area of Bandwidth-on-Demand (BoD) a new set of issues or does it (re-)arrange an old group of known solutions? Activities-on-Demand (AoD) solved the batch processing problems some years ago. BoD did provide an early support of data storage problems in high-performance computing. Today, the need for communication bandwidth at different locations, at different times, and with different qualities is termed BoD. Computing-on-Demand (CoD) solved lately on-demand number cruncher problems for, *e.g.*, very large simulations, and provided short-term computing needs in a quite static manner. Devices-on-Demand (DoD)[5] serves today for a multitude of desktops, laptops, palmtops, and mobile phone needs of the human and his/her selection choice according to his/her communication, entertainment, and computing demands. Energy-on-Demand (EoD) essentially targets at the minimization of energy demands, while keeping up the current service delivery requirements (cf. George Huitema's talk). Fiber-on-Demand (FoD) describes the lambda switching approaches developed. Resources-on-Demand (RoD) mechanisms are deployed in Grids and they address the related CPU, memory, and library resources assignments in distributed systems. Services-on-Demand (SoD) methods deliver dedicated services, such as simulations, specified forecasts, and banking services, while the similar abbreviation may encompass in a much broader sense Software-on-Demand (SoD). This does cover today's Software-as-a-Service (SaaS) approach in terms of samples like text processing, logistics, or similar solutions for many commercial software products. Finally, XoD equals the term "X-on-Demand", where X $\in$ "requested for item of interest". Based on the discussion within the seminar's plenary, XoD may be termed "*oD" (pronounced as "Star-on-Demand") as well, which simply follows in this case a different syntactical agreement.

Listing that many sample approaches being worked on in the past outlines one aspect, but in which way do they differ? In which way do they compare to each other? Thus, commonalities between all of those XoDs include the object

[5] In the traditional networking and infrastructure work, DoD is often referred to as the Department of Defense, which sponsored the initial development of Internet-like protocols and the basic Internet architecture, but did not address any on-demand aspects.

and its uses, while the demand may be known by quality, by quantity, and by time, in turn a clear result of this demand upon request will be (a) possible – including the guarantee –, while accepting the request and (b) not possible – while rejecting the request. If the demand is not known, respective mechanisms should be in place as well; they will map the specifically determined demand at run-time onto available system components under the assumption that some results can be achieved. Of course, in this case guarantees cannot be granted at all. In both cases the definition of the "requested item of interest" is as important as the operation methods for those items themselves. In addition, the mode of operation may be either centralized, like for brokers, coordinators, or within a bandwidth exchange point, or it may run in a (fully) decentralized manner, which covers an ATM-type (Asynchronous Transfer Mode) of style, where ATM switches operate independent of each other, or a P2P-type (Peer-to-peer) of style, in which independent nodes collaborate independently and without any centralized guideline with other peers.

The four key dimensions to address in such a XoD setting include the following ones. This is especially essential for defining relevant resources, their use, and the field of interactions, in which they are located in.

1. Timing – In terms of duration, date, month, minutes, seconds, hours, milliseconds, and the like, the on-demand has to be specified and quantified.
2. Resource Allocation Mechanisms – In combination with the resource description-to-user view optimizations may be possible in technical aspects or in economics, which themselves may be applied in a pre-determined scheme or on-line.
3. Resource Allocations – They determine the resource description-to-technology view and its technical schemes in terms of – technology-dependent – lambdas, protocols, queues, ports, and the like, which the research domain does know about today quite well.
4. Interacting Roles – ISP-to-ISP (Internet Service Provider) interactions define a very realistic case for on-demand approaches in many on those Xs, while the ISP-to-user interactions may determine a more futuristic case, in which more than 64 Kbps bandwidth chunks can be assigned to end-users in a residential environment.

Thus, a selected group of open issues (at least) needs to address the following questions, each of them in the dedicated context of the "requested for item of interest":

1. Regulation – Which roles need to be determined by law? Who deals with whom? Is non-repudiation a must for on-demand services?
2. Business Models – Policies for interacting domains are essential for a common interpretation of technical parameters. Charging models for ISPs and their respective incentive models are necessary for a successful cooperation.
3. Missing Mechanisms – To enable an XoD in many of those cases outlines above, the respective technical layer Y signaling protocol for technology Z has to be developed, which requires an integrated authentication and authorization mechanism.

Finally, driven by these observations described above, the re-visiting of commonalities has to be performed. Can the research community ensure that generalizations of these mechanisms determined above are possible? Many of these allocation algorithms, roles in consideration, and signaling principles in use will be applicable for different Xs, while the concrete definition of the scheme, the instantiated role, or the protocol implemented may differ due to technology constraints. If such commonalities may not be identified, the discussion on "on-demand" approaches may address unnecessary problems and solutions developed may re-invent the wheel in different forms over and over again. Therefore, as of today the BoD approaches undertaken do re-visit a set of known solutions as such (reservations, allocations, mechanisms, and optimizations), sometimes improving those, sometimes integrating a similar type of inherent problem. However, they target newer technology, which leads to new signaling protocols, algorithmic instances, and additional performance gains compared to existing work. Thus, BoD – interpreted in the networking domain specifically – does address new challenges, which include those four dimensions outlined above and which need to answer questions on regulations and economics at the same time.

Keywords: X-on-Demand

Self-managed Inter-domain Pricing: An Overview and Discussion of Possible Approaches

Bruno Tuffin (IRISA/INRIA Rennes, FR)

Telecommunication networks are now an interconnection of competitive operators that need to cooperate to deliver and ensure end-to-end delivery of traffic. Inter-domain agreements have to be performed, and pricing is seen as a relevant way to reward intermediate domains for forwarding the traffic of others, so that the network can be efficiently self-managed. In the devoted literature, Vickrey-Clark-Groves (VCG) auctions have been highly considered because they provide proper incentives, lead to an efficient use of the network and verify other relevant characteristics. On the other hand, it has been highlighted in this specific context that they are not budget-balanced, nor robust to collusion. After reviewing the list of properties that we ideally would like the inter-domain pricing to verify and the fact that no mechanism can check them all together, we discuss the properties that could be relaxed. We present all the related works not especially considering inter-domain pricing, and discuss how they apply and what the corresponding results for this specific problem are.

Joint work of: Tuffin, Bruno; Tran Hoang, Hai;

Resource Allocation and Provision for Bandwidth on Demand in SINET3

Shigeo Urushidani (National Institute of Informatics – Tokyo, JP)

In this talk, I presented bandwidth on demand capabilities in the new Japanese academic network, called SINET3. First, main service and network features of the network were presented. After the detailed explanation of the networking technologies deployed in the network, the general architecture for BoD services was presented. I clarified the service parameters of BoD services and some points on path calculation by taking into account deployed layer-1 technologies. I also showed sample reservation screens for users and actual interfaces between layer-1 BoD server, developed by NII, and layer-1 devices. I also showed current projects using layer-1 BoD services, eVLBI, high-quality remote backup, and t-room projects, and real setup and release times for these services. Finally, I showed open issues in our network.

Bandwidth on Demand Contract(s)

Martin Waldburger (Universität Zürich, CH)

With the ongoing commercialization of electronically provided value-added services in the Internet, the need for commercialization support mechanisms has become apparent. In addition to a reliable authentication, authorization, and accounting of service and resource usage, legal compliance in automated contract formation determines a critical issue for successful commercialization. Accordingly, an in-depth understanding of the respective contract nature is to be obtained in order to reflect the respective desired legal outcome consistently during contract negotiation, performance, and the potential contract enforcement phase in case of disputes.

Bandwidth on Demand contracts constitute a valuable study object for a contract qualification in terms of a business-to-business application case for a two-party international contract construct under private law. In particular, main and secondary obligations of both, service provider and service customer, need to be assessed – a task whose outcome is highly dependent on those terms and guarantees to be included in the according SLA. This assessment leads in a second step to an analysis of available options to bundle these identified and characterized main and secondary obligations into a single (probably quite complex) contract construct or into separate (potentially easier to be managed) contracts, each covering a single identified main obligation only.