

Resilient Integration of Distributed High-Performance Zones into the BelWue Network Using OpenFlow

Michael Menth*, Mark Schmidt*, Daniel Reutter*, Robert Finze†, Sebastian Neuner‡, and Tim Kleefass‡

* Chair of Communication Networks, University of Tuebingen, Tuebingen, Germany

† Zentrum fuer Datenverarbeitung, University of Tuebingen, Tuebingen, Germany

‡ BelWue Koordination, University of Stuttgart, Stuttgart, Germany

Abstract—BelWue is the Internet service provider for higher education and research institutions in Baden-Wuerttemberg, Germany. Recently, high-performance zones (HPZs) have been established on major university campuses and interconnected with a high-speed network for innovation and research (NeIF). This work presents the SDN-NeIF architecture, a resilient integration of the HPZs into the NeIF and the legacy infrastructure of BelWue and its connected universities, leveraging OpenFlow and BGP. The concept is validated by a prototype, results from a field trial are provided, and additional benefits of using software-defined networking (SDN) in this context are discussed.

I. INTRODUCTION

The Internet service provider (ISP) BelWue interconnects 45 higher education and research institutions in Baden-Wuerttemberg, Germany, including 9 university campuses, altogether about 150 locations. There is a trend towards service centralization among the universities in Baden-Wuerttemberg, i.e., some data- or computation-intensive services are offered only by single institutions and are available to others only via the BelWue network. Therefore, a network for innovation and research (NeIF, Netzwerk fuer Innovation und Forschung) has recently been set up to interconnect the university campuses through a flexible, optical network with 100 Gb/s wavelengths between neighboring sites, divided into 10×10 Gb/s bandwidths. It enables 10 Gb/s point-to-point connections to provide low-latency services between any two campuses but their overall number is limited by the switching matrices at each site. Therefore, a full optical mesh among all campuses is not feasible. As the existing campus infrastructure of universities is not ready to support these high data rates and the optical network of BelWue is already on its way to be upgraded to multiple 100 Gb/s, so-called high-performance zones (HPZs) with high-performance hosts (HPHs) are established and directly connected to the NeIF via OpenFlow-capable, Ethernet-based border switches (BSs).

The NeIF is already partly used to provide point-to-point connections for special demands to carry data-intensive traffic between university locations. In contrast, the HPZs are currently operated in isolation to the existing (legacy) infrastructure and can be interconnected through point-to-point

connections provided by the NeIF (see Figure 1). In this context, the project bwNET100G+ was set up among research groups and the computation centers at the Karlsruhe Institute of Technology, the University of Tuebingen, the University of Ulm, and the BelWue coordination. Its objective is to make networking within and among universities more flexible leveraging novel networking technologies like OpenFlow, and to improve transport layer and security aspects to enable university users to benefit from the increased bandwidths in the BelWue network.

The work presented in this paper is an outcome of the bwNET100G+ project. It suggests the SDN-NeIF architecture, a resilient integration of the HPZs into the NeIF and the legacy infrastructure of BelWue and its connected universities. It uses OpenFlow technology and BGP for that purpose, and does not require additional hardware. The use of software-defined networking (SDN) is attractive because it makes networking more flexible and allows for improved security, traffic engineering, and more cost-efficiency. This work can be seen as one part in the pre-planing for an evolution towards an SDN-enabled next-generation ISP platform of BelWue.

The rest of the paper is structured as follows. Section II discusses similar activities connecting special zones within and among universities with high-speed networks. Section III presents the conceptual integration of the HPZs into the NeIF and the existing infrastructure of BelWue and the universities leveraging OpenFlow and BGP. Section IV describes the implementation of a prototype and Section V reports results from a field trial. Section VI discusses opportunities of the SDN-based HPZ integration. Finally, Section VII concludes this work.

II. RELATED WORK

We briefly review similar projects that facilitate high-speed communication for scientific applications within or between university campuses. Some of them leverage SDN technology.

McCahill [1] proposed at an Internet2 technical meeting a high-speed bypass around a university's core network for special traffic, e.g., scientific data. Departments are connected to the core network with SDN-capable switches. They are interconnected through dedicated high-speed links and decide which traffic to bypass over the high-speed network. The goal is to provide high bandwidth between different departments or

This work has been supported in the bwNET100G+ project by the Ministry of Science, Research and the Arts Baden-Wuerttemberg (MWK). The authors alone are responsible for the content of this paper.

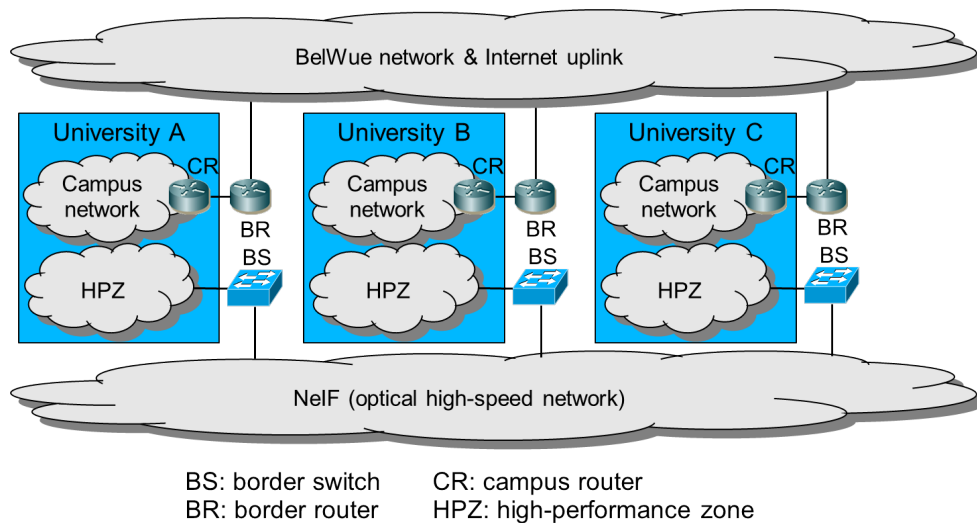


Fig. 1. The HPZs belong to the universities. BelWue interconnects them through a flexible, optical network (NeIF). The HPZs are currently operated in isolation from the production infrastructure.

different locations of the same university for scientific data and to relieve the campus core network.

The ESnet Science DMZ [2], [3] defines high-speed zones within universities that are separate from the campus network. Within a university, there may be multiple DMZs with high-performance equipment, e.g., computation and storage servers, and with special security policies and enforcement, e.g., for different projects. The DMZs are used for high-performance scientific applications. They are connected via the university's border router to the campus network and the Internet. SDN technology is used for communication within and among the DMZs of a single university. High-speed connections at 100 Gb/s among the DMZs of a single university enable the use of the resources in different DMZs outside the campus network so that large data volumes do not need to be relayed through the campus network. There is no special high-speed network directly interconnecting the Science DMZs of different universities.

SciPass [4] describes a security-enhanced Science DMZ. It uses an intrusion detection system (IDS) to classify traffic as trusted and untrusted. Trusted traffic is, e.g., scientific data that is exchanged between different universities. SciPass uses OpenFlow switches as load balancers for IDS. After a flow is classified as trusted, the network is configured so that this flow can bypass firewalls and is routed around potential bottlenecks. The goal of this approach is to facilitate the utilization of a 100 Gb/s inter-campus connectivity.

Internet2 [5] is a network connecting US education and research centers. It is based on an optical 100 Gb/s backbone and reaches from the US West coast to the US East coast. The typical uplink of a location is 10 Gb/s. The network can be used, e.g., to interconnect Science DMZs at different locations. The more general goal of Internet2 is to provide a high-speed network for collaborative applications, distributed research experiments, as well as grid-based data computation and analytics.

III. RESILIENT SDN-BASED INTEGRATION OF HPZS

As illustrated in Figure 1, the HPZs are connected to BSs that are interconnected through a flexible, optical platform – the NeIF. However, the use of the NeIF in this context is undefined so far and the HPZs are operated in isolation, but can be connected via the NeIF for research purposes. There is no connection to the production environment and the Internet.

In the following, we discuss requirements for the integration of HPZs into the NeIF and the legacy infrastructure. The integration is achieved through the BSs which are equipped with appropriate forwarding rules. BGP is used to make the HPZs reachable from the campus networks and the Internet. Resilience mechanisms effect that traffic is rerouted via the BelWue core network if the NeIF fails. Finally, we discuss required information exchange between BelWue and universities.

A. Requirements

As the HPZs belong to different institutions, their equipment belongs to different IP number spaces so that the entirety of all HPZs cannot be operated as a single layer-2 network. Furthermore, the BSs of the HPZs must be managed only by BelWue while the equipment within the HPZs is managed by the universities. The HPZs must be reachable from university campuses and the Internet, and the reachability information must be communicated in an automated way. The interconnection of the HPZs should require only a few optical links because the ability of establishing optical point-to-point links in the NeIF should be retained for special applications. Therefore, only BSs between neighboring sites are connected through single-hop optical links and BSs relay data among HPZs using packet switching. While university campuses are redundantly connected to the existing BelWue network via two border routers (BRs) and disjoint paths, which is omitted in Figure 1, the NeIF currently exhibits a tree structure. Therefore, the resilience of the communication among HPZs against link failures should also benefit from their integration into the legacy BelWue network.

B. Interconnection of Border Switches

A /22 IPv4 and an IPv6 address space are reserved for the entirety of all HPZs out of which each HPZ receives its own /24 prefix. As illustrated in Figure 2, any HPZ is connected to the NeIF platform through an OpenFlow-capable BS which is operated by BelWue and controlled by a BS controller (BSC) run by BelWue. The BS has 10 Gb/s interfaces and a direct link to the campus router (CR) which is typically located in the university's data center. The BS should have a dedicated connection to the BR so that BR and BS can exchange traffic without forwarding it through the campus network. The BS either directly connects devices within an HPZ or may talk to gateways which hide from the BS the remaining network structure within the HPZ. The BSs have an optical link towards the neighboring BSs in the NeIF. If a single 10 Gb/s channel does not suffice to carry the traffic between neighboring BSs in the NeIF, several 10 Gb/s channels may be bonded using the Link Aggregation Control Protocol (LACP).

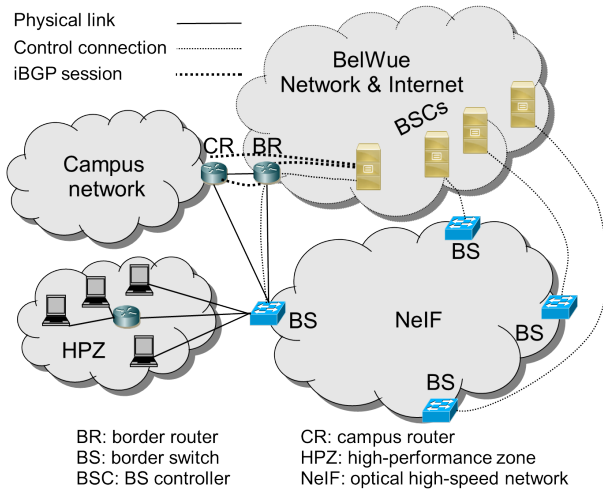


Fig. 2. An OpenFlow-capable border switch (BS) attaches an HPZ to the campus network, the legacy BelWue infrastructure, and to other HPZs over a network of OpenFlow switches.

C. Forwarding Rules for the BS

The BS requires one static rule to reach its BSC via the BR. In addition, the BS needs rules to act as a gateway for the HPZ. For directly connected devices dynamically generated forwarding rules are used. The addition of these rules works as follows. A match rule for the /24 HPZ prefix is configured on the BS. It has low priority and triggers the export of the packet header to the BSC. The BSC sends an encapsulated ARP request to the BS which then broadcasts this request to its neighbors and returns the result to the BSC. Then, the BSC installs a new forwarding rule with higher priority for the requested device on the BS. To avoid ARP requests being broadcast into other HPZs, an additional rule on the BS blocks all non-encapsulated ARP requests received from the NeIF. If a gateway within the HPZ is connected to the BS, the BS requires a rule to forward traffic towards the prefix behind the gateway. This rule can either be configured statically or with the help of a routing protocol.

The BSC also installs appropriate forwarding rules for the IP prefixes of the other HPZs on the BS. Gateways and directly

connected devices have the BS configured as default gateway, e.g., by DHCP. Moreover, the CR and the BR see the BS as potential next IP hop. However, if the BS forwards traffic to a neighboring BS in the NeIF, it does not change the source and destination MAC address because packets are forwarded by the BSs only by their destination IP addresses. When forwarding traffic into the HPZ, to the CR, or to the BR, the BS does set appropriate MAC addresses.

D. HPZ Reachability via BGP

The reachability of the HPZ from the campus and the Internet is achieved through BGP. As the BS is only a switch, it cannot speak BGP itself and requires that the BSC acts as proxy to maintain BGP connections with the BR and the CR.¹ The BSC announces the reachability of the /22 prefix of the entirety of all HPZs and the /24 prefix of the local HPZ to the BR and CR with the BS being the next hop. The BR propagates this information further to the Internet and in particular to other BRs. The CR announces itself as the next hop for the university campus to the BSC and BR. Therefore, the BSC installs a rule on the BS to forward traffic destined to the campus to the CR. The BR announces a default route to the CR and the BSC so that it becomes the next hop for the CR and the BS for traffic to the Internet.

E. Resilience Mechanisms

For resilience purposes, the university campus is connected via two CRs with identical IP address to the BR. The Virtual Router Redundancy Protocol (VRRP) [6] is used between them so that they act as one virtual router which ensures connectivity even if one of them fails. In a similar way, the university uplink is realized via two BRs with identical IP address over disjoint paths to the BelWue core network, and here is a full mesh interconnection among the two BRs and CRs. The figures omit this complexity for the sake of clarity. Because of this arrangement, any BR, CR, or path towards the BelWue core network may fail without compromising the uplink of a location. As a result, the connection between the BS and its BSC can be considered as highly reliable because the BSC is located in the BelWue core network.

Within the NeIF, the BSs locally detect if a link fails between them and inform their BSCs about this event. In such a case, a BSC withdraws the /22 HPZ prefix via BGP to the CR and the BR. If the link failure is repaired, the BSC is notified and reannounces the /22 HPZ prefix. In case of a failure, the CR and the BR still forward traffic for the local HPZ to the BS, but they forward traffic for all other HPZs via the BR and the BelWue core network to the BRs of the corresponding HPZs. Traffic towards the local HPZ is rerouted by the other BS detecting the failure via its BR, the BelWue core network, the local BR, and the local BS. If that traffic originates in another HPZ, it loads the uplink of a location which is actually not involved in the communication. Therefore, we currently work on controller-to-controller communication so that the BSC can inform other BSCs whose BSs are no longer reachable through the NeIF to reconfigure their BSs and to withdraw the /22

¹This functionality is already supported by many controllers, but they can act as proxy only for a single node. Therefore, our design requires a separate BSC for every BS.

HPZ prefix to their BRs and CRs. As a result, affected traffic towards a local HPZ will be rerouted at its origin. We further work on controller resilience which is not yet covered by our current prototype.

F. Information Exchange between BelWue and Universities

In the presented architecture, BelWue controls the BR, the BSs, the BSCs, and the NeIF. The universities control the equipment in the HPZs and the campus networks including the CRs. Therefore, some information needs to be exchanged between BelWue and universities with attached HPZs.

Devices in the HPZ that are directly attached to the BS respond to ARP requests from the BS so that the BSC can dynamically configure appropriate forwarding rules on the BS. Possibly, this can be simplified, e.g., by an automatic export of ARP mappings from a DHCP server within the HPZ to the BSC. Then, the BSC can install forwarding rules for all devices directly connected to the BS so that the above described mechanism is not needed to set up dynamic forwarding rules.

The BS needs to be configured with the prefixes which are reachable through attached gateways in the HPZ. The university may communicate this information to BelWue so that the BSC can install appropriate forwarding rules on the BS. As an alternative, the attached gateways may communicate this information through routing protocols via the BS to the BSC in an automated way.

The presented integration makes only a few assumptions about the connection of campus networks and HPZs to BelWue’s infrastructure which are generally met. The only change to the existing campus network is the addition of one BGP neighbor to the CR. All other changes are restricted to the BelWue infrastructure.

IV. PROTOTYPE IMPLEMENTATION

We first implemented the concept for the HPZ integration on Mininet and in a local testbed [7]. Then, we implemented a prototype on the target platform. As the BelWue and campus networks are production environments, they must not be used for experiments. Therefore, we use only the NeIF and the BSs as physical components and virtualize most other components of the architecture (BRs, campus hosts (CHs), high-performance hosts (HPHs) in the HPZ, and an Internet host (IH)) as virtual machines (VMs) on servers. Thereby, the prototype can run on the target platform while being isolated from the production infrastructure. In the following, we explain the mapping of physical and virtual components to experimental hardware and briefly describe the virtualization platform.

A. Prototype Design

The HPZs are currently equipped with a test rack. It contains the access to the NeIF, a management VPN, a management switch, an HP ProCurve 3500 switch with 1 Gb/s interfaces, an HP ProCurve 5406 switch compatible to OpenFlow version 1.3 with 10 Gb/s interfaces, and 5 servers with 10 Gb/s interfaces. We utilize the testbed equipment of the HPZs in Tuebingen and Ulm. Access to this equipment is realized via VPN to the management switch which has direct links to

an additional network interface of the equipment. This has the advantage that management traffic does not influence the experiments in the testbed and that access to the components is possible even in case of an error or misconfiguration.

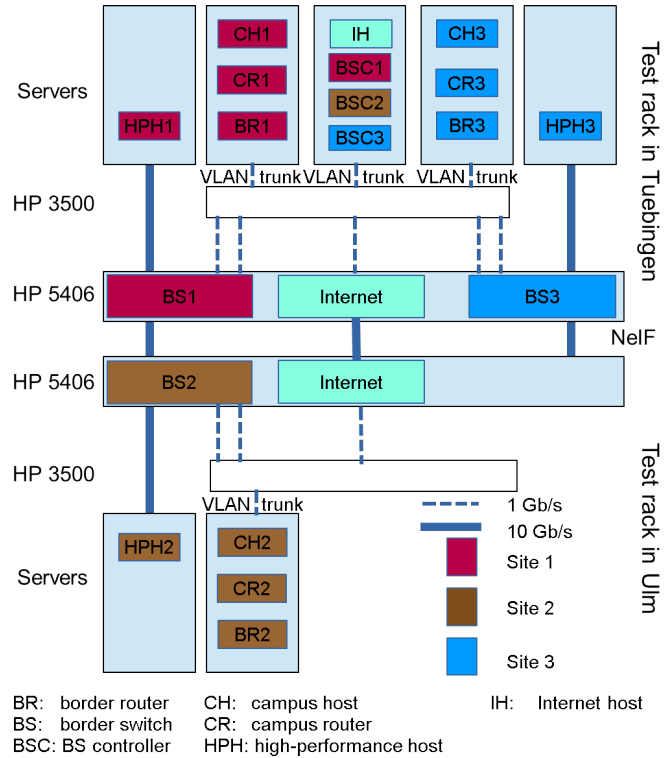


Fig. 3. Mapping of physical and virtualized network nodes to experimental hardware.

The prototype is illustrated in Figure 3 which omits the management network. It represents three different sites, each consisting of a virtualized campus network and a HPZ that are interconnected via BSs through the NeIF. The entities in the figure have site-specific colors. The two HP 5406 switches are subdivided into two and three partitions, respectively. Three partitions are used for the BSs and another two for the Internet connection between the three sites. BS1 is connected via the NeIF to BS2, and BS2 is connected via the NeIF to BS3. This is the central part of the prototype which runs on physical machines. One HPH per HPZ is implemented as a VM on a dedicated server and connected to the corresponding BS. The BR and CR also run as VMs on a different server and have also a direct link to the BS. Some servers host several VMs but have only a single port. Therefore, traffic from the VMs is carried in different VLANs as a trunk between the servers and the HP 3500 switch. The HP 3500 switch connects the CR to the CH and the BR, it interconnects the BS with the BR, and it facilitates communication between the BR and the corresponding BSC, as well as among BR1, BR3, and the IH. The virtualized BRs and CRs are based on Quagga and the BSC leverages the Ryu platform. All IP addresses are statically configured so that a DHCP server is not needed and debugging is simplified. IP addresses are also directly configured on BSCs so that host discovery is not needed.

B. Server Virtualization Platform

The servers are equipped with two Intel Xeon E5 processors, 128 GB RAM, three SSDs which are assembled to a RAID5, and two Intel 10 Gb/s NICs, one for management purposes and one for experiments. Each server can host several VMs. As virtualization platform we use KVM as hypervisor in conjunction with qemu. All VMs and hosts run Ubuntu Linux as operating system with some basic tools for debugging and performance analysis. In [7] we describe how the virtualization techniques are used to build a local testbed for SDN-NeIF. The virtualization concept for the prototype is almost the same. The HP ProCurve 3500 series switches de-/multiplex the VLANs of different VMs of a server to/from different physical switch ports. As a result, the BS and the CR are connected only with untagged VLAN. This is even more realistic because components connect the BSs on dedicated switch ports. Furthermore, the BS is not required to handle VLANs which keeps the BSC simple.

V. FIELD TRIAL

The prototype in Figure 3 corresponds to the logical experiment setup in Figure 4. This field trial involves the real NeIF and physical BSs while all other network entities are virtualized substitutes for production nodes. We performed basic connectivity tests, measured TCP throughput between selected nodes, and checked failover behavior for selected link failures. Additional tests are ongoing.

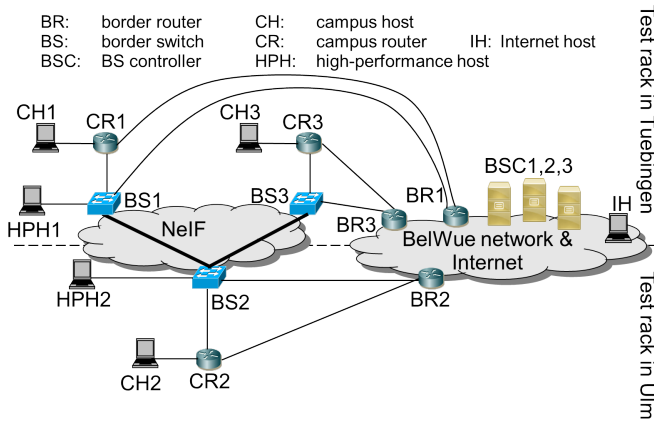


Fig. 4. Logical structure of the prototype in Figure 3.

A. Basic Connectivity Tests

We used the *ping* and *traceroute* utilities for reachability tests. We verified that

- CH1 reaches IH via CR1 and BR1.
- CH1 reaches HPH1 via CR1 and BS1.
- CH1 reaches HPH3 via CR1, BS1, the NeIF, and BS3.
- CH1 reaches CH3 via CR1, BR1, BR3, and CR3.
- HPH1 reaches CH1 via BS1 and CR1.
- HPH1 reaches IH via BS1 and BR1.

- HPH1 reaches CH3 via BS1, BR1, BR3, and CR3.
- HPH1 reaches HPH3 via BS1, the NeIF, and BS3.

A local roundtrip time between HPH1 and CH1 takes 0.2 ms, a roundtrip time over a single NeIF link from HPH1 to HPH2 takes 2.1 ms, and a roundtrip time over two NeIF links from HPH1 to HPH3 takes 4.05 ms.

B. Throughput Tests

We measured throughput from a single TCP connection between selected nodes using *iperf*. HPH1 and HPH3 are connected in the prototype through a path with 10 Gb/s links and we observed a throughput between them of 9.4 Gb/s. HPH1 and CH1 are connected through a path with a 1 Gb/s bottleneck link and we measured a traffic rate of up to 960 Mb/s between them. As the traffic between CH1 and CH3 is forwarded via BR1 and BR3, it traverses twice the 1 Gb/s links between the server and the HP3500 switch in the prototype. Therefore, we could only achieve a traffic rate of 480 Mb/s. Thus, we could leverage almost the entire link capacity with only a single TCP connection. This concludes that the implementation of the prototype is rather efficient.

C. Failover Tests

We validated the rerouting for link failures in the NeIF using the *ping* and *traceroute* utility. If the link between BS1 and BS2 fails, traffic from HPH2 to HPH1 is rerouted by BS2 via BR2, and BR1 to BS1 and HPH1. Thus, the legacy uplink of the originating location protects the failure. Traffic from HPH3 to HPH1 is first forwarded via BS3 to BS2 which then also reroutes it via BR2, and BR1 to BS1 and HPH1. Hence, the legacy uplink of an intermediate location protects against the failure. As mentioned before, this may be avoided through controller-to-controller communication. If the link failure is repaired, the normal forwarding behavior is restored.

The described recovery process requires that the failure-detecting BS notifies its BSC which then sends a withdraw to the BR and the CR for the /22 prefix. This approach is only slightly slower compared to the use of a BGP router instead of the BS because the BGP router can immediately withdraw the /22 prefix after failure detection.

VI. OPPORTUNITIES OF SDN-BASED HPZ INTEGRATION

The proposed integration of HPZs is simple, cost-efficient, and resilient. Apart from that, the SDN-based architecture easily allows to carry campus-to-campus traffic over the NeIF, facilitates improved security and traffic engineering, and offers the perspective on alternative redundant uplinks. Moreover, it can be incrementally deployed which may be an important step towards a cost-efficient SDN-based wide area network (WAN). We discuss these issues in the following.

A. Carrying Campus-to-Campus Traffic over the NeIF

The NeIF offers very high transmission capacities that could be leveraged to carry campus-to-campus traffic within the BelWue at a speed of up to 100 Gb/s. This may facilitate the usage of centralized data-intensive services provided by

particular university computation centers. Examples are storage and computation clusters.

Carrying campus-to-campus traffic requires only reconfiguration of the BS through the BSC to send traffic destined to other university campuses through the NeIF and to announce the other campus via BGP to the CR. As the number of university campuses within the BelWue is low, configuring the additionally required forwarding rules on the BS is not problematic.

B. Improved Security

The use of SDN technology offers flexibility that may be used for improved security in high-speed networks. Currently, we work on OpenFlow-assisted firewall offloading for high-speed networks and on security concepts for the use of resources in remote HPZs. SciPass [4] leveraged OpenFlow switches to automatically detect scientific data flows that may be safely bypassed around IDS systems. This concept may be reused by bwNET100G+.

C. Improved Traffic Engineering

With SDN, more flexible traffic management can be supported than with conventional routing. This feature can be leveraged if the CR forwards all outbound traffic to the BS. The BS may be configured by its BSC to forward to the NeIF only traffic from certain applications or traffic between certain departments or project groups. It is also possible to carry only selected flows, e.g., elephant flows, through the NeIF while all other flows use the legacy BelWue network via the BR. However, per-flow forwarding may require significantly more forwarding rules on the BS.

Large data transfers between remote HPZs may be automatically scheduled to transmit them at high data rates and to avoid impact on other traffic during busy hours. Such an approach has already been taken by Google [8].

Traffic from the HPZ forwarded via the BR into the legacy BelWue network may overload the existing infrastructure and cause quality of service (QoS) degradation for traffic from the campus network. To avoid that, rate limiting on the link from the BS to the BR may be applied. This feature is available from OpenFlow version 1.3 onwards. Such rate limitations may be applied possibly only to traffic from other universities which possibly emerges in case of failures and rerouting.

D. Alternative Redundant Uplink

A university campus typically has a redundant uplink. To that end, it is connected with two BRs and two physically disjoint paths to the legacy BelWue network. With the help of the BS and the NeIF it is possible to provide a redundant uplink with only one BR and the BS at every location. To that end, the CR should forward all traffic to the BS and the BS forwards desired traffic to the BR. If the BR or the uplink fails, the BSCs need to be notified in some way and reconfigure BRs through iBGP and the BSs such that that affected traffic is carried through the NeIF. Possibly, the traffic can be load-balanced. Requiring only one BR at each location saves operational costs and acquisition costs if a BR has to be replaced.

E. Cost-Efficient SDN-Based WAN

There is always a run in upgrading to the next magnitude of bandwidth. An IP router usually needs to be replaced as whole. This is costly in an environment like the BelWue with many sites but only a few institutions per site. Therefore, it is attractive to use the available high bandwidth on the optical layer in a clever way while saving expenses for high-cost devices. An SDN-based WAN is a vision for the future, but is difficult to achieve in practice. ISPs are rather reluctant to introducing SDN technology in their networks because they must assure stable operation. Their administrators require some time to get familiar with the new control paradigm and develop appropriate debugging tools. The presented approach can be incrementally deployed. First, HPZs may be just interconnected via the NeIF. Then, communication with the Internet may be facilitated. Later, resilience may be added. Afterwards, the discussed advanced features may be introduced. Incremental deployment offers the possibility to integrate a test network with initially non-critical applications, and using it for production purposes only if sufficient test and operation experience has been gained. Thus, the presented concept simplifies the move towards an SDN-based WAN.

VII. CONCLUSION

The ISP BelWue interconnects university campuses via a legacy network and their HPZs through a high-speed optical network, the NeIF. We presented a resilient, OpenFlow-based integration of the HPZs into the NeIF, other existing BelWue infrastructure, and the university campus networks. The proposed SDN-NeIF architecture is designed such that it can be fully controlled by BelWue but gives enough flexibility to cooperating universities. It does not need additional hardware, it is resilient and scalable. Furthermore, it facilitates improved security and traffic engineering, simplifies a redundant uplink for attached universities, and can be incrementally deployed. The latter is important for the move towards an SDN-based WAN which is attractive for cost-efficiency. The implementation of a prototype and the reported field trial are first steps in that direction.

REFERENCES

- [1] Mark McCahill, "SDN, IDM, and Research Computing at Duke," <http://meetings.internet2.edu/media/medialibrary/2015/10/19/20151007-McCahill-SDN-IDM-ResearchComputing-Duke.pdf>, accessed: 2016-07-20.
- [2] E. Dart, L. Rotman, B. Tierney, M. Hester, and J. Zurawski, "The Science DMZ: A Network Design Pattern for Data-intensive Science," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, 2013.
- [3] ESnet, "Science DMZ," <https://fasterdata.es.net/science-dmz/science-dmz-architecture>, accessed: 2016-07-20.
- [4] GlobalNOC, "SciPass," <https://globalnoc.iu.edu/sdn/scipass.html>, accessed: 2016-07-20.
- [5] The Internet2 Community, "Internet2," <http://www.internet2.edu>, accessed: 2016-07-20.
- [6] S. Nadas, "RFC5789: Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6," March 2010.
- [7] M. Schmidt, R. Finze, D. Reutter, and M. Menth, "Demo: Resilient Integration of Distributed High-Performance Zones into the BelWue Network Using OpenFlow," in *International Teletraffic Congress*, Würzburg, Germany, Sep. 2016.

- [8] S. Jain, A. Kumar, S. Mandal, J. Ong, A. Poutievski, Leon Arjun Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, “B4: Experience with a Globally-Deployed Software Defined WAN,” in *ACM SIGCOMM*, Hong Kong, China, Aug. 2013.