

Similarly, introducing virtualization and wide-area delays will increase the *variability* of the service provided by the network function, due to issues such as network congestion, network failures or oversubscribed clouds.

Finally, network functions may be exposed to unusual threats from their co-tenants in a cloud. This may require further hardening of services, requiring more stringent or additional firewall protection, adding to the cost and delay in NFV.

4. CONCLUSIONS AND FUTURE WORK

This paper has focused on the *deployment* issues encountered as part of designing, deploying, and evaluating a transparent Split TCP middlebox for wireless networks. There are other issues we encountered but, due to space limitations, cannot include them here. Broadly, they fall under two categories:

Debugging. These issues centered on development decisions that aided or hindered debugging running systems, particularly under high load, or after extended deployments (i.e., months).

Measurement. While it was simple to measure performance improvements in a dedicated testbed, it was much more difficult to quantify to customers how well the system worked for *their* traffic and environment. In particular, *variability* introduced by using real workloads, networks and servers made this much more challenging. Constraints required by the customers also restricted our actions.

We are currently writing up our entire experience including the above in longer form and hope to publish that as a full paper.

Acknowledgements

We would like to thank Hani Jamjoom and the anonymous reviewers for their helpful comments. This research was sponsored by US Army Research laboratory and the UK Ministry of Defence and was accomplished under Agreement No. W911NF-06-3-0001. Vasilis Pappas is currently at Google. His contribution to this work was performed while at IBM.

5. REFERENCES

- [1] A. Bakre and B. Badrinath. I-TCP: Indirect TCP for mobile hosts. In *IEEE ICDCS*, 1995.
- [2] Baracuda. VLAN deployments. <https://techlib.barracuda.com/bwf/deplyvlan>.
- [3] V. Bernat. Coping with the TCP TIME-WAIT state on busy linux servers. <http://vincent.bernat.im/en/blog/2014-tcp-time-wait-state-linux.html>.
- [4] Blue Coat. Technology primer: VLAN tagging. https://www.bluecoat.com/sites/default/files/documents/files/VLAN_Tagging.1.pdf.
- [5] J. Border, M. Kojo, J. Griner, G. Montero, and Z. Shelby. Performance enhancing proxies intended to mitigate link-related degradations. *IETF RFC 3135*, 2001.
- [6] K. Brown and S. Singh. M-TCP: TCP for mobile cellular networks. *ACM SIGCOMM CCR*, 27(5):19–43, 1997.
- [7] Cisco Systems. Cisco 8500 series wireless controllers. <http://www.cisco.com/c/en/us/products/wireless/8500-series-wireless-controllers/index.html>.
- [8] Cisco Systems. Enterprise network architecture v. 3. http://www.cisco.com/web/strategy/education/us_education/borderless_education.html.
- [9] Cisco Systems. Overlay Transport Virtualization (OTV). <http://www.cisco.com/c/en/us/solutions/data-center-virtualization/overlay-transport-virtualization-otv/index.html>.
- [10] Cisco Systems. WCCP version 2. <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipapp/configuration/x3sg/iap-xe-3sg-book/iap-wccp-v2.html>.
- [11] ETSI. Network functions virtualization: Architectural framework. http://www.etsi.org/deliver/etsi_gs/nfv/001_099/002/01.01.01_60/gs_nfv002v010101p.pdf, 2013.
- [12] M. Honda, Y. Nishida, C. Raiciu, A. Greenhalgh, M. Handley, and H. Tokuda. Is it still possible to extend TCP? In *Proc. of ACM IMC*, 2011.
- [13] T. Kivinen, B. Swander, A. Huttunen, and V. Volpe. Negotiation of NAT-Traversal in the IKE. *IETF RFC 3947*, 2005.
- [14] Linux Kernel Sources. tproxy. <http://www.kernel.org/doc/Documentation/networking/tproxy.txt>.
- [15] M. Luckie and B. Stasiewicz. Measuring path MTU discovery behaviour. In *Proc. of ACM IMC*, 2010.
- [16] M. Mahalingam, D. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridhar, M. Bursell, and C. Wright. VXLAN: A framework for overlaying virtualized layer 2 networks over layer 3 networks. *IETF RFC 7348*, 2014.
- [17] M. Mathis and J. Heffner. Packetization layer path MTU discovery. *IETF RFC 4821*, 2007.
- [18] J. Mogul. Network behavior of a busy Web server and its clients. Technical Report 95/5, DEC WRL, Oct. 1995.
- [19] B. Pfaff, J. Pettit, K. Amidon, M. Casado, T. Koponen, and S. Shenker. Extending networking into the virtualization layer. In *Hotnets*, 2009.
- [20] J. Postel. Transmission control protocol. *IETF RFC 793*, 1981.
- [21] V. Sekar, N. Egi, S. Ratnasamy, M. K. Reiter, and G. Shi. Design and implementation of a consolidated middlebox architecture. In *Proceedings of the 2012 USENIX NSDI*.
- [22] Serverfault. Why would a server not send a SYN/ACK packet in response to a SYN packet. <http://serverfault.com/questions/235965/why-would-a-server-not-send-a-syn-ack-packet-in-response-to-a-syn-packet>. Accessed: 2015-04-01.
- [23] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else's problem: Network processing as a cloud service. In *Proceedings of the ACM SIGCOMM 2012*.
- [24] Squid Wiki. SquidFAQ interception proxy. <http://wiki.squid-cache.org/SquidFaq/InterceptionProxy>.
- [25] R. Yavatkar and N. Bhagawat. Improving end-to-end performance of TCP over mobile internetworks. In *WMCSA*, pages 146–152. IEEE, 1994.