Traffic Monitoring Considered Reasonable

Mark Allman
International Computer Science Institute

I. Position

Networking and security researchers must grapple with the following four potential constraints in the course of their work: (i) legal issues, (ii) institutional policies, (iii) community norms and (iv) personal ethics. These possible constraints can be intertwined and inform one another. In this position paper we set aside all except the third constraint: community norms. We do not mean to diminish the importance of considering the others, but, rather aim to make the case that the community has organically developed a community norm with respect to traffic monitoring.

Position: *network traffic monitoring—broadly defined—fits well within the networking and security research community's norms*. We believe this position is demonstrated through the community's rough consensus and running code.

II. Community History

The following illustrative references show that our community has leveraged network monitoring (i) for over two decades, (ii) for a variety of purposes, (iii) using various types of monitors—from packet traces to application layer logs, with many stops in-between—and (iv) across a large number of investigators and institutions.

[1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30], [31], [32], [31], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47], [48], [49], [50], [51], [52], [53], [54], [55], [56], [57], [58], [59], [60], [61], [62], [63], [64], [65], [66]

Further, an even larger number of referees, program committee members and editorial boards have considered and accepted this work, which is an indication of the broad level of acceptance traffic monitoring enjoys across the community.

III. Reasoning

The above illustrative history is not enough to justify our position. Rather, as discussed in the Menlo Report [67] we must also consider the benefits and harms of network traffic monitoring. Therefore, in addition to the historical record we also offer several additional points:

- There is little doubt that the *benefit* from understanding the reality of large-scale network and security phenomena through in-situ observation has been immense. It is self-evident that if we do not monitor networks we will have only a theoretical understanding of how they operate and this would create a large divide between research and reality in both the networking and security domains. However, this point is not only self-evident, but we also know from history that we as a community have learned much from empirical studies of Internet behavior across many dimensions (i.e., see the illustrative history given in § II).
- While the benefits of network monitoring are clear, we must also consider the possible *harms* caused by such endeavors. When the harm becomes significant we should consider an activity at least dubious.
- When considering the potential harms of traffic monitoring, the list of possibilities is nearly endless—largely because traffic is ultimately triggered by human activity. The potential harms that could come from gathering and analyzing network traffic range from personal privacy issues to institutional embarrassment to business implications—each of which has an attendant list of consequences.
- While we can conjure all manner of *potential harm* that stems from monitoring networks, cases where network monitoring has lead to *actual harm* are quite rare. We believe this is largely because of the care taken by the research community when monitoring operational networks. We discuss care of network data further in § IV.
- Additionally, we note that researchers are often not studying users, per se. This further reduces the potential for harm because users do not have to be identified and their activity understood as part of our analysis (we address this point further in § IV).

Given the above consideration of the benefits, harms and history of network monitoring we believe *the presumption should be that investigations that observe it-situ activities should be considered well within the bounds of what the community considers to be ethically reasonable*. This presumption does not excuse all behavior, but we believe it should be the starting point.

IV. Care

The community's history of using empirical Internet observation is obvious. However, "because we have always done it" is not a good reason for continuing a given practice. Rather, we note that the history gives us an understanding of the benefit and harm but also an idea about how to conduct careful network monitoring that minimizes potential harm ("running

code"). Before distilling a set of principles for exercising care with respect to Internet data, we offer three pieces of context:

- Observations of an operational network ultimately are quite often observations about real peoples' actions. Often we are not studying particular human behavior, but rather some facet of the system (e.g., use of DNS fast flux or round-trip time assessments). However, we should never lose track of the fact that traffic data often comes from human beings and therefore there is always the potential for harm that stems from observing specific activities.
- While traffic is often triggered by human behavior, we also should remember that traffic is not stamped with the identity of the human being(s) involved. That does not mean it is impossible to determine who is involved in the recorded traffic. We can sometimes pick up a small bit of information that clues us in to who was involved in some specific traffic. Further, we have mechanisms to link various bits of traffic together (e.g., addresses, cookies, referrer information). Taken together we can sometimes use small breadcrumbs to piece together a broader story about a user's network usage. However, we note that this does not happen without effort. In other words, the data we collect is not naturally of the form "here is what Alice did yesterday on the Internet ...". This is an important distinction because in the majority of the cases there is no reason for researchers to construct these sorts of user profiles and therefore user privacy is naturally obscured.
- Finally, note that the potential sensitivity also depends on the specific dataset and vantage point. For instance, fine-grained packet traces taken within a backbone network are generally less sensitive than the same measurement within a department in a University. First, the scale of the former is such that digging out information about specific people is more difficult. Second, while a department-level trace likely includes a comprehensive view of an individual's activity, only part of that activity is likely to hit any particular backbone. While our contention is that the community has viewed both these cases as reasonable, they illustrate that there is likely no one-size-fits-all approach to thinking about traffic observations.

Our intention is not to enumerate a specific and fine-grained set of best practices about careful data handling. Rather, we offer three high-order principles researchers should think about as they collect, archive and use network traffic data.

- **Logistics:** Researchers should exercise care in terms of what data is collected, how the data is archived and who has access to the data. A particularly useful strategy is leveraging the "need to know" principle. In other words, data should not be available to someone unless they have a direct need to access the data. Further, data should be provided in the granularity required to assess specific questions. For instance, a packet trace that includes payloads of all traffic observed may be winnowed to only packet headers involving TCP port 80 before being given to a student to perform a study on HTTP. In addition, machines holding sensitive data should be well secured—including proper access control, application of relevant security updates and monitoring for possible breakins.
- **Aggregate:** When reporting analysis of network traffic data users should not be identified. Fortunately, we are generally trying to convey insights across a breadth of traffic (and hence users) and therefore aggregating across users is a natural approach which also obscures individual users' activity. On occasion it is useful to show the behavior of a single user as an anecdote to better describe some phenomenon. This can be accomplished without identifying the user. And, further can often be presented in a way that elides the time the measurement was taken and the specific endpoints of the communication. In other words, by focusing on the underlying phenomenon, even an anecdote that is directly tied to a specific user can be presented in a way that cannot be traced back to the given user.
- **No Fishing:** Analysis of data should have a purpose. Researchers should focus on a question—even if only an inkling—and not simply surf through collected data to see what they might see.

Ultimately, exercising care with Internet traffic data requires researchers to carefully *think* about how they are conducting their measurements and analysis. While we believe the community has largely demonstrated an ability to conduct traffic monitoring without causing harm, we must remain vigilant about how we go about such activities.

V. Broader Points

We close with three broad points:

- Our position is about traffic monitoring—at various vantage points—and not about all network and security measurement. While we believe there is rough consensus that traffic monitoring is both useful and causes little direct harm, we do not believe that all empirical research activities share the same consensus (e.g., BotNet infiltration).
- Our intention is to sketch what we believe is an organically developed community norm. Even if we are correct in positing this norm, it is not the last word on the appropriateness of a given experiment. As noted above, researchers must show care and the required care depends on the specific data and vantage point.
- As we note in § I, researchers face additional constraints in the form of laws, institutional policies and personal ethics. Therefore, even a general and widely agreed upon community norm can dictate behavior in every situation.
- While there are indeed many open ethical dilemmas surrounding network and security measurement, it may be useful to start writing down areas where there is (rough) consensus on what is reasonable.

REFERENCES

[1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic," in *ACM SIGCOMM*, 1993.

[2] K. Claffy, H. Braun, and G. Polyzos, "Traffic Characteristics of the T1 NSFNET Backbone," in *IEEE INFOCOM*, 1993.

[3] V. Paxson, "Empirically-Derived Analytic Models of Wide-Area TCP Connections," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316–336, Aug. 1994.

[4] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, Jun. 1995.

[5] M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," in *ACM SIGMETRICS*, May 1996, pp. 160–169.

[6] M. Arlitt and C. Williamson, "Web Server Workload Characterization: The Search for Invariants (Extended Version)," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, Oct. 1997.

[7] W. Lee, S. Stolfo, and K. Mok, "Mining Audit Data to Build Intrusion Detection Models," in *International Conference on Knowledge Discovery and Data Mining*, Aug. 1998.

[8] B. Huffaker, J. Jung, D. Wessels, and K. Claffy, "Visualization of the Growth and Topology of the NLANR Caching Hierarchy," in *WWW Caching Workshop*, Jun. 1998.

[9] J. Bennett, C. Partridge, and N. Shectman, "Packet Reordering is Not Pathological Network Behavior," *IEEE/ACM Transactions on Networking*, Dec. 1999.

[10] P. Barford and M. Crovella, "Measuring Web Performance in the Wide Area," *Performance Evaluation Review: Special Issue on Network Traffic Measurement and Workload Characterization*, Aug. 1999.

[11] Y. Zhang and V. Paxson, "Detecting Stepping Stones," in *USENIX Security Symposium*, Aug. 2000.

[12] A. Sridharan, S. Bhattacharyya, C. Diot, R. Guerin, J. Jetcheva, and N. Taft, "On the Impact of Aggregation on the Performance of Traffic Aware Routing," in *International Teletraffic Congress*, 2001.

[13] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS Performance and the Effectiveness of Caching," *Networking, IEEE/ACM Transactions on*, vol. 10, no. 5, pp. 589–603, 2002.

[14] N. Brownlee, K. Claffy, and E. Nemeth, "DNS measurements at a root server," in *IEEE Global Telecommunications Conference (GLOBECOM)*, 2002.

[15] M. Allman, E. Blanton, and W. Eddy, "A Scalable System for Sharing Internet Measurements," in *Passive and Active Measurement Workshop*, Mar. 2002.

[16] M. Bykova and S. Ostermann, "Statistical Analysis of Malformed Packets and Their Origins in the Modern Internet," in *ACM SIGCOMM/USENIX Internet Measurement Workshop*, 2002.

[17] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman, "A Measurement-Based Analysis of Multihoming," in *ACM SIGCOMM*, 2003.

[18] M. Roughan, S. Sen, O. Spatscheck, , and N. Duffield, "Class-of-Service Mapping for QoS: A Statistical Signature-based Approach to IP Traffic Classification," in *ACM SIGCOMM/USENIX Internet Measurement Conference*, 2004.

[19] J. Pang, A. Akella, A. Shaikh, B. Krishnamurthy, and S. Seshan, "On the Responsiveness of DNS-based Network Control," in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*. ACM, 2004.

[20] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, "Characteristics of Internet Background Radiation," in *ACM Internet Measurement Conference*, Oct. 2004.

[21] H. Dreger, A. Feldmann, V. Paxson, and R. Sommer, "Operational Experiences with High-Volume Network Intrusion Detection," in *ACM CCS*, Oct. 2004.

[22] S. Uhlig, "High-order Scaling and Non-stationarity in TCP Flow Arrivals: a Methodological Analysis," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, Apr. 2004.

[23] J. Jung, V. Paxson, A. Berger, and H. Balakrishnan, "Fast Portscan Detection Using Sequential Hypothesis Testing," in *IEEE Symposium on Security and Privacy*, May 2004.

[24] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "BLINC: Multilevel Traffic Classification in the Dark," in *ACM SIGCOMM*, 2005, pp. 229–240.

[25] C. Park, F. Hernandez-Campos, S. Marron, and F. Smith, "Long-Range Dependence in a Changing Internet Traffic Mix," *Computer Networks*, vol. 48, no. 3, Jun. 2005.

[26] C. Kreibich, A. Warfield, J. Crowcroft, S. Hand, and I. Pratt, "Using Packet Symmetry to Curtail Malicious Traffic," in *ACM HotNets*, 2005.

[27] N. Duffield, C. Lund, and M. Thorup, "Learn More, Sample Less: Control of Volume and Variance in Network Measurement," *IEEE Transactions in Information Theory*, vol. 51, no. 5, 2005.

[28] E. Cooke, Z. M. Mao, and F. Jahanian, "Hotspots: The Root Causes of Non-Uniformity in Self-Propagating Malware," in *Proc. of DSN*, 2006.

[29] G. Liang, N. Taft, and B. Yu, "A Fast Lightweight Approach to Origin-Destination IP Traffic Estimation Using Partial Measurements," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2634–2648, june 2006.

[30] K. Cho, K. Fukuda, H. Esaki, and A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic," in *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 4. ACM, 2006, pp. 207–218.

[31] D. Moore, C. Shannon, D. J. Brown, G. M. Voelker, and S. Savage, "Inferring Internet Denial-of-Service Activity," *ACM Trans. Comput. Syst.*, vol. 24, no. 2, pp. 115–139, May 2006.

[32] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *SIGCOMM*, 2006.

[33] J. C. Mogul and M. Arlitt, "SC2D: An Alternative to Trace Anonymization," in *Proc. ACM MineNet Workshop*, 2006.

[34] E. Kohler, J. Li, V. Paxson, and S. Shenker, "Observed Structure of Addresses in IP Traffic," *ACM/IEEE Transactions on Networking*, vol. 14, no. 6, Dec. 2006.

[35] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan, "Analyzing Large DDoS Attacks using Multiple Data Sources," in *ACM SIGCOMM Workshop on Large-Scale Attack Defense*, 2006.

[36] G. Bartlett, J. Heidemann, and C. Papadopoulos, "Inherent Behaviors for On-line Detection of Peer-to-Peer File Sharing," in *IEEE Global Internet*, May 2007.

[37] C. Reis, S. Gribble, T. Kohno, and N. Weaver, "Detecting In-Flight Page Changes with Web Tripwires," in *Proc. NSDI*, 2008.

[38] G. Thatte, U. Mitra, and J. Heidemann, "Detection of Low-Rate Attacks in Computer Networks," in *IEEE Global Internet Symposium*, Apr. 2008.

[39] A. Schulman, D. Levin, and N. Spring, "On the Fidelity of 802.11 Packet Traces," in *Passive and Active Measurement Conference*, Apr. 2008.

[40] A. Mahimkar, J. Yates, Y. Zhang, A. Shaikh, J. Wang, Z. Ge, and C. T. Ee, "Troubleshooting Chronic Conditions in Large IP Networks," in *ACM CoNext*, Dec. 2008.

[41] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching Television Over an IP Network," in *ACM/USENIX Internet Measurement Conference*, Oct. 2008.

[42] S. Hao, N. A. Syed, N. Feamster, A. G. Gray, and S. Krasser, "Detecting spammers with SNARE: Spatio-temporal network-level automatic reputation engine," in *Usenix Security Symp.*, 2009.

[43] A. Medem, M.-I. Akodjenou, and R. Teixeira, "A. Medem, M.-I. Akodjenou, and R. TeixeiraTroubleMiner: Mining Network Trouble Tickets," in *IFIP/IEEE International Workshop on Management of the Future Internet*, Jun. 2009.

[44] G. Maier, A. Feldmann, V. Paxson, and M. Allman, "On Dominant Characteristics of Residential Broadband Internet Traffic," in *ACM Internet Measurement Conference*, Nov. 2009.

[45] T. Benson, A. Akella, and D. A. Maltz, "Mining Policies From Enterprise Network Configuration," in *ACM Internet Measurement Conference*, Nov. 2009.

[46] M. Canini, W. Li, M. Zadnik, and A. W. Moore, "Experience with High-Speed Automated Application-Identification for Network-Management," in *Symposium on Architectures for Networking and Communications Systems*, Oct. 2009.

[47] F. Qian, A. Gerber, Z. M. Mao, S. Sen, O. Spatscheck, and W. Willinger, "TCP Revisited: A Fresh Look at TCP in the Wild," in *ACM/USENIX Internet Measurement Conference*, Nov. 2009.

[48] H. Yin, X. Liu, F. Qiu, N. Xia, C. Lin, H. Zhang, V. Sekar, and G. Min, "Inside the Bird's Nest: Measurements of Large-Scale Live VoD from the 2008 Olympics," in *ACM/USENIX Interent Measurement Conference*, Nov. 2009.

[49] M. Afanasyev, T. Chen, G. Voelker, and A. Snoeren, "Usage Patterns in an Urban WiFi Network," *IEEE/ACM Transactions on Networking*, vol. 18, no. 5, Oct. 2010.

[50] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin, "A First Look at Traffic on Smartphones," in *ACM SIGCOMM/USENIX Internet Measurement Conference*, 2010.
[51] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Houston, "Internet Background Radiation Revisited," in *ACM SIGCOMM/USENIX Internet Measurement Conference*, Nov. 2010.
[52] V. Sekar, M. K. Reiter, and H. Zhang, "Revisiting the Case for a Minimalist Approach for Network Flow Monitoring," in *ACM/USENIX Internet Measurement Conference*, Nov. 2010.
[53] A. Gember, A. Anand, and A. Akella, "A Comparative Study of Handheld and Non-Handheld Traffic in Campus WiFi Networks," in *Passive and Active Measurement Conference*, 2011.
[54] S. Coull, F. Monrose, and M. Bailey, "On Measuring the Similarity of Network Hosts: Pitfalls, New Metrics, and Empirical Analyses," in *Network & Distributed System Security Symposium*, Feb. 2011.
[55] G. Stringhini, T. Holz, B. Stone-Gross, C. Kruegel, and G. Vigna, "BotMagnifier: Locating Spambots on the Internet," in *USENIX Security Symposium*, Aug. 2011.
[56] H. H. Song, Z. Ge, A. Mahimkar, J. Wang, J. Yates, , and Y. Zhang, "Analyzing IPTV Set-Top-Box Crashes," in *ACM SIGCOMM Workshop on Home Networks*, Aug. 2011.
[57] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. Voelker, and S. Savage, "Show Me the Money: Characterizing Spam-advertised Revenue," in *USENIX Security Symposium*, Aug. 2011.
[58] F. Schneider, B. Ager, G. Maier, A. Feldmann, and S. Uhlig, "Pitfalls in HTTP Traffic Measurements and Analysis," in *Passive and Active Measurement Conference*, 2012.
[59] D. McCoy, A. Pitsillidis, G. Jordan, N. Weaver, C. Kreibich, B. Krebs, G. Voelker, S. Savage, and K. Levchenko, "PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs," in *USENIX Security Symposium*, 2012.
[60] Y. Xie, F. Yu, Q. Ke, M. Abadi, E. Gillum, K. Vitaldevaria, J. Walter, J. Huang, and Z. M. Mao, "Innocent by Association: Early Recognition of Legitimate Users," in *ACM CCS*, 2012.
[61] V. Paxson, "Internet Traffic Archive," http://ita.ee.lbl.gov/.
[62] "A Community Resource for Archiving Wireless Data At Dartmouth (CRAWDAD)," http://crawdad.cs.dartmouth.edu.
[63] "Internet Measurement Data Catalog (DatCat)," http://www.datcat.org/.
[64] "Protected Repository for the Defense of Infrastructure Against Cyber Threats (PREDICT)," http://www.predict.org/.
[65] "LBNL/ICSI Enterprise Tracing Project," http://www.icir.org/enterprise-tracing/.
[66] "Click Dataset," http://cnets.indiana.edu/groups/nan/webtraffic/click-dataset.
[67] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan, "The Menlo Report," *IEEE Security and Privacy*, vol. 10, no. 2, 2012.