

New Frontiers in Internet Network Management

Ehab Al-Shaer⁺, Albert Greenberg[†], Charles Kalmanek[‡], David A. Maltz[†],
T. S. Eugene Ng[§], Geoffrey G. Xie^{*}

⁺University North Carolina Charlotte, [†]Microsoft Research, [‡]AT&T Labs Research, [§]Rice University,
^{*}Naval Postgraduate School

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

Abstract

Network management represents an architectural gap in today's Internet [1]. Many problems with computer networks today, such as faults, misconfiguration, performance degradation, etc., are due to insufficient support for network management, and the problem takes on additional dimensions with the emerging programmable router paradigm. The Internet Network Management Workshop is working to build a community of researchers interested in solving the challenges of network management via a combination of bottoms-up analysis of data from existing networks and a top-down design of new architectures and approaches driven by that data. This editorial sets out some of the research challenges we see facing network management, and calls for participation in working to solve them.

Categories and Subject Descriptors: C.2.3 [Computer-Communication Network]: Network Operations

General Terms: Management

Keywords: Network & service management, Internet Network Management Workshop, research questions

1. INTRODUCTION

In many ways, computer network management remains the least understood aspect of computer networking. There is a lack of well-established principles guiding the design of networks for manageability. There is also a lack of scientific understanding of the evolution of network state in real operational environments. Commercially, given the huge dollar values involved in managing networks, a number of companies have created products to help network operators gain visibility over and manage the behavior of their networks. However, none of these products come even close to addressing the challenges, and their users remain deeply unsatisfied with the results, despite significant investment to develop and deploy such products.

The difficulties of network management range from the mundane to the failure of some of the best “organizing principles” of networking. For example, network devices are too often still managed one box at a time rather than managed as an integral networked system, despite the fact that the correct operation of a network requires related updates to two devices to be either both accepted or both rejected with a transaction-like semantic. Underlying this somewhat

mundane example is a more fundamental challenge: ensuring the consistency of the vast network state. As an organizing principle, “soft state” has proven to be extremely valuable for achieving the eventual consistency of protocol state in many individual network protocols. However, given the vast amount of state that needs to be managed in each network device, relying on the soft state principle alone may not be enough.

The Internet Network Management Workshop (INM) is a relatively young venue created to foster state-of-the-art research in network management. INM differs from other venues in that its focus is on tackling fundamental network management problems with a combination of a bottom-up analysis of data from existing network systems and a top-down design of new systems grounded in the theories derived from the bottom-up analysis. The INM workshop seeks to elevate participants' collective experience with operational IP networks into concepts, principles, and theories that can be leveraged in today's networks and carried forward into clean-slate designs that intrinsically support management rather than treating management as a bolted-on after-thought.

In this editorial, we articulate some of the research issues that motivated us to start the INM workshop and discuss a few possible research directions.

2. RESEARCH CHALLENGES

New Environments, New Challenges.

These are exciting times for network management research as new programmable router paradigms are dramatically increasing the potential for control over the network elements. Platforms such as OpenFlow [9] and software routers [2, 6, 4, 7, 11] create the need for new network abstractions and mechanisms to manage these systems and tie them into secure networks that operate with high reliability. While existing networking systems have drawbacks, many bugs have been driven from implementations over decades of use and refinement. Time frames on that scale are not affordable for new systems. Since all new systems have new bugs, the community must challenge itself to create new classes of debugging tools that are compatible with and take the full advantage of the new abstractions. Some of the new platforms host the control plane in a small number of servers: we also need algorithms for improving the network's ability to survive the misbehavior of these critical components.

The emerging cloud computing and data center infrastructures face increasing user expectations for reliability and performance, yet they present significant network and service management challenges as a result of their scale and complexity. In these settings, the sheer number of servers and network elements under management are leading to new and fresh ideas for large scale systems management [5, 3], protecting the integrity of computation and data even in the face of numerous failure and unpredictable workloads. INM is interested in the role that networking should play in making data centers economical, which includes issues of agility (e.g., giving each data center tenant the illusion it has its own private expandable data center or data centers), resource usage (e.g., ensuring each tenant obtains the performance and resources it needs), and manageability (e.g., ensuring that few humans are able to operate very large data centers while minimizing down time).

In addition, the emergence of overlay networks as a virtual network on the top of the Internet introduces new types of operational and management problems such as multi-level failures propagation due to interactions between overlay and underlay, and interdomain policy violation due to bypassing service providers rules. Unlike native IP networks, overlays have dynamic connectivity that makes monitoring, correlating and diagnosing these problems much harder.

Scalable Yet Rigorous State Management.

Much of network management is centered on managing the distributed state that are necessary to implement a service or to monitor the health of the network. Therefore, research ought to take a serious look at the challenges involved in maintaining the consistency of network state. Traditionally, network protocols have relied heavily on the notion of periodically refreshed soft-state to achieve and maintain eventual state consistency. However, with the state explosion and accompanying increase in network service complexity, making every piece of state soft may not be feasible. One potential direction is to consider other points in the space of consistency models, from looser partial consistency to stronger atomic, transactional consistency semantics. Even in situations where these concepts seem particularly helpful and natural, such as network configuration, change management, and on-line network debugging, they have not been well-explored.

Management Friendly Protocols and Data-Plane Primitives.

There is already a huge number of existing mechanisms used in various aspects of network management. There is an acronym soup of routing, signaling, QoS and virtualization protocols; however, they must be knitted together to provide network services, and this is where innovation is needed. To reduce the complexity of network management we argue for characterizing the behavior of these existing mechanisms, and then creating useful abstractions for them. INM is interested in new methods for design and modeling of control plane protocols and data-plane primitives. Relevant questions include: What prerequisites are required by a mechanism? What invariants does a mechanism maintain? What performance guarantees does a mechanism make? What metric to use to assess both the performance and cost (e.g., complexity) of such a mechanism? What are the possible failure modes of a mechanism? What

are the supported methods of recovery? Answers to these questions will inform network management systems to maximize potential automation in service provisioning, network monitoring, failure diagnostic, and failure recovery.

Going further, future mechanisms could also be designed to explicitly support interactions and coordination with other mechanisms. Most of the existing mechanisms are designed for a specific isolated purpose. They rarely expose any programmatic interfaces. Yet, in practice, multiple mechanisms often need to work in conjunction to realize the network's objectives. As a result, today's networks often resort to customized glue logic and hacks to integrate mechanisms that do not have explicit interfaces for interactions. One such example is using route-redistribution for gluing together multiple routing protocols [8]. Configuration hacks like these are responsible for many outages, as they substantially increase the complexity of the system and make it much harder to automate network management. Therefore, in designing new mechanisms, it is worth considering what explicit programmatic interfaces should be supported to allow component integration to be seamless. An added advantage of defining these interfaces is the potential to enhance the network performance via joint optimization of the parameter settings of multiple mechanisms. For example, more efficient reachability control may be possible by jointly optimizing the configurations of packet filters and routing protocols [10]. Such interface should play a key role to achieve a balance between multiple trade-offs (i.e., constraints) in accomplishing the goal such as intrusiveness vs. accuracy, usability vs. risk, that might be infeasible to achieve based on a view of a single system.

Testbeds, Data, and Evaluation Methods.

One barrier to improving the impact and quality of network management research has been the lack of publicly available, minimally sanitized network configuration data, traffic traces, and operational experience data (e.g., outage and error information). INM solicits research tools for creating, sanitizing, sharing these types of data broadly among the research community. By working together, we are finding that our individual successes at obtaining access to data and configuration can be leveraged into greater benefits for all.

Another factor slowing down the progress of network management research is the primitive state of the scientific methods available for studying network management problems and for evaluating solutions. We need both formal methods for analyzing network management systems and equally important, a set of benchmarks including performance metrics and complexity metrics for comparing network management solutions. Experimental evaluation is also closely tied to data availability, as compelling benchmarks need to be based on realistic workflows, real trace derived tests, realistic reconfiguration tasks, realistic offered workload, etc. Environments like Emulab (and GENI in the future) provide the beginnings of testbeds for network management, but more needs to be done on simulator support for network management research, as well as realistic testbeds for repeatable emulation. A mature set of accessible scientific methodologies will be a great catalyst for accelerating innovative research in network management.

Leveraging Both Static and Dynamic Verification.

Modern network systems are highly complex and prone to bugs and other vulnerabilities. Software bugs could cause, for example, a router to mis-process packets. Software vulnerabilities could be exploited to allow an attacker to inject misbehaviors that undermine the network's functions. Therefore, even if the network management system correctly manages a network's state under normal conditions, the runtime behavior of the network may be degraded or incorrect under adverse conditions. It is therefore useful to begin thinking about how to architect a network so that it has built-in support for operational correctness verification. The challenging issues to address are: What configuration models can support global and scalable static verification analysis? How to compare and rank different network configurations? How to instrument, observe and debug the network in timely and synchronized manner? How should data and control plane behavior be reliably and securely monitored and reported? How could a network management system automatically derive the reference correct behavior at device level? How can a massive amount of monitored data be processed to verify correctness?

3. PARTICIPATION

We are looking for others who want to tackle these challenges and join the conversation. The next INM workshop will be co-located with NSDI 2010 in San Jose, CA. INM is seeking case studies, experimental results, position papers, as well as provocative ideas and clean-slate designs. Paper registrations are due on November 30, and the submission deadline is December 7. The full call for papers is at <http://www.usenix.org/events/inm10>.

4. REFERENCES

- [1] V. Cerf, B. Davie, A. Greenberg, S. Landau, and D. Sincoskie. FIND Observer Panel Report April, 2009.
- [2] M. Dobrescu, N. Egi, K. Argyraki, B.-G. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh, and S. Ratnasamy. RouteBricks: Exploiting parallelism to scale software routers. In *SOSP*, 2009.
- [3] S. Ghemawat, H. Gobiuff, and S.-T. Leung. The Google File System. In *SOSP*, 2003.
- [4] M. Handley, O. Hudson, and E. Kohler. XORP: An open platform for network research. In *HotNets*, October 2002.
- [5] M. Isard. Autopilot: Automatic data center management. *Operating Systems Review*, 41(2), 2007.
- [6] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Trans. Computer Systems*, August 2000.
- [7] T. V. Lakshman, T. Nandagopal, R. Ramjee, K. Sabnani, and T. Woo. The SoftRouter architecture. In *HotNets*, November 2004.
- [8] F. Le, G. G. Xie, D. Pei, J. Wang, and H. Zhang. Shedding light on the glue logic of the internet routing architecture. In *SIGCOMM*, pages 39–50, 2008.
- [9] Openflow. <http://www.openflowswitch.org/>.
- [10] Y.-W. E. Sung, S. G. Rao, G. G. Xie, and D. A. Maltz. Towards systematic design of enterprise networks. In *CoNEXT*, 2008.
- [11] H. Yan, D. A. Maltz, T. S. E. Ng, H. Gogineni, H. Zhang, and Z. Cai. Tesseract: A 4D network control plane. In *NSDI*, April 2007.