Human-Network Interaction: the Weakest Link?

Rudra Dutta, Computer Science, North Carolina State University
Whitepaper submitted to “Looking Beyond the Internet: SDI/SDX Workshop”

As computing, networking, storage technologies converge, and grow increasingly more sophisticated, and increasingly automated into agile software, an often-overlooked vulnerability in these information systems is gaining in risk even faster: the management, administration, and operation of such systems. The threat posed by misconfiguration of computers or networking systems have always been recognized, but are seldom the subject of active information system engineering research. What research exists has largely paid attention to the configuration and management of individual end systems, because it has been reasonably assumed by many researchers that these are the only systems that lay persons will configure – surely networks and datacenters, large or small, are configured by professionals, who are presumably highly qualified and can be depended upon to do the job of configuration and administration without error; after all, a comparatively simple job.

Unfortunately this is far from the truth. Recent surveys by our research group confirm that most network administrators (and even those styled “engineers” and “architects”) in enterprise networks lack advanced degrees in computing fields, or even an undergraduate degree. Indeed, many lack a college degree of any kind – or a college degree from unrelated fields in humanities or basic science. Many are also paid very little, get very little on-the-job training, and lack a clear growth path within the profession. In light of these findings, it is not surprising to see incidents such as the ones in 2004, when BGP misconfiguration in network AS9121 (TTnet - the largest ISP in Turkey) resulted in misdirected or lost traffic for tens of thousands of networks and downtime of more than 10 hours [Alin2005, Nanog23]; in 2012, when 40% of Google's users experienced near complete loss of service for 18 minutes due to a combination of errors some of which were human errors [Ars Technica, 2012/12]; or when nearly 5000 domains, including Yelp, LinkedIn, Fidelity and the United States Postal Service, were taken down in the infamous June 2013 Ztomy DNS misconfiguration – ironically, the intended configuration action had been intended to be a defensive response to a perceived DDoS attack [Schultz2013, Cisco blog].

It may be assumed that more automation will solve this problem, or reduce it to negligible proportions, but experience teaches us otherwise. As the configuration management of simple protocols and systems is handed over to automation by ever-more intelligent (therefore, complex) protocols (usually to eliminate repetitive and tedious tasks, and/or manage growing scale), it is certainly true that (a) the overall amount of configuration workload reduces, *but at the same time* (b) the complexity and conceptual background required increases, *and* (c) the consequences of misconfiguration grow more dire and far-reaching. A simple example is provided by IP forwarding tables. In ARPANET days, they were agreed upon over the phone and at in-person meetings, and hand-configured into routers – that too by experts. To release experts from such tasks, dynamic routing protocols were postulated, which in turn allowed the network to scale by many orders of magnitude, and eventually enable more advanced functionality. Instead of configuring hundreds (and by now, what would be between thousands and hundreds of thousands) of forwarding entries by hand, only a few OSPF or BGP configurations have to be set. But the forwarding table entries were transparent, the IP protocol comparatively simple to understand, and a consequence of misconfiguration in one would be localized – whereas OSPF2/3 or BGP4, even though they hide most of their complexity from the configuration interface, demand deeper understanding of far more complex protocols to configure correctly (or worse, are blindly configured by rote, as arcane knowledge), and errors in one place can bring down large sections of the Internet e.g. the 2010 Duke U experiment that brought a Cisco bug to light, and in the process disrupted over 3500 route prefixes in 60 countries in the minutes it took human operators to stop the experiment [IDG News Service, Aug 2010].

The recently popular vision of SDN-based systems as a unifying future architecture of network-enabled computing systems has been seen by some as essentially eliminating the configuration issue. While there may be eventual truth to this, it is also true that configuration will now simply move to a more abstract, and more powerful, plane. As the job of the underpaid, undertrained network operations employee changes from configuring by hand to writing controller app code that can execute thousands of time in the seconds taken for a human operator to realize there is an error, we can look forward to ever more spectacular failures.

At NCSU, we believe that such problems must be attacked at the root – by (a) understanding human errors, and being able to model them, (b) detecting human errors at the moment when they are made, not from failures afterward, (c) using technology to correct or prevent human errors, ideally before the configuration goes live. We have been working on studying the human process of network administration and management, both by observing and modeling it. We have published some of our observations in [Mushi, Dutta, 2014 Workshop on Human Centered Big Data Research] and [Mushi, Murphy-Hill, Dutta, DRCN 2015]. We have surveyed network architects and administrators from a number of enterprises, and we have also created a virtual environment to study the configuration workflow in detail – at this moment, about 50 graduate student volunteers are exercising this environment, and in the near future we expect to distribute it to partner organizations.

Further, we see the SDN paradigm as providing a way to interpose sanity-checking code between the human and the configurations. Our work in flow specific service-insertion using OXM in OpenFlow networks [Udechukwu, Dutta, ICNP-CoolSDN 2014] may be used to allow such services to run at arbitrary locations in the network. We have demonstrated the practicality of such a system at a demonstration using commercial equipment at the 1st Global Cities Teams Challenge, Washington DC, June 2015 (in association with Extreme Networks). We commend this area as a critical one deserving further attention by a broad spectrum of researchers as we look beyond the Internet as it is today.