Measuring the 4:11 Effect:

The Power Failure and the Internet

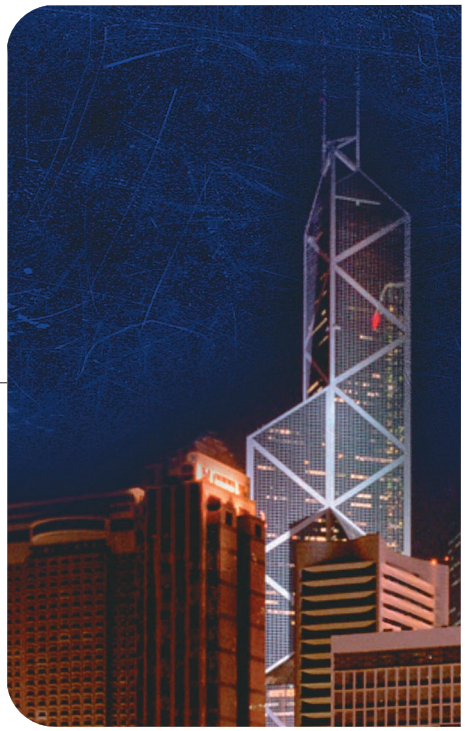

If you were one of the unfortunate millions affected by the Northeast blackout of 2003, then loss of Internet connectivity probably was not your biggest problem. Whether trapped in a subway, stranded in an airport, simmering in unrelenting heat, or just sitting in the dark, 50 million people experienced inconveniences far more trying than loss of email and Web surfing capacities.

Dennis McGrath
Dartmouth College

But for those of us who think about critical infrastructures and cascading infrastructure failures, the blackout that began at 4:11PM Eastern Daylight Time (EDT) on 14 August provides insight into the overlap between the power grid and the information grid. Just as 9/11 was a wakeup call for national security, 4:11 was a wakeup call for the defenders of critical infrastructures, and an opportunity to learn more about the their complex interdependencies.

For years, we've speculated a lot about the Internet's potential threat to the power grid, but considerably less about a vulnerable power infrastructure's threat to the Internet. Common sense tells us that all of our infrastructures (including transportation, banking, telecommunication, water supply, and so on) depend on each other—and particularly on the power grid—but we don't understand this symbiotic relationship very well. Understanding begins with measurement and quantification, so we must measure the power failure's effect on the Internet. For details on the measurement process, see the "How do we measure it?" sidebar.

The blackout

Internet watchers were on high alert during the second week of August 2003. It had been nearly three weeks since disclosure of the Microsoft Distributed Component Object Module (DCOM) vulnerability, and security experts predicted a worm's imminent appearance that could make previous worms look like child's play. Just six months earlier, the Slammer worm took nearly everyone by surprise, and no one wanted to be caught off guard again. So, in early August, we carefully watched the Internet health indicators for any signs of worm activity. Border Gateway Protocol (BGP) watchers were particularly on edge, because experience showed that scanning worms have a pronounced effect on BGP activity.[1]

When the highly anticipated Blaster worm finally struck on the afternoon of 11 August, Internet performance metrics showed that its effects were less than expected. Not that Blaster was harmless, but unlike previous worm epidemics, which caused large amounts of Internet disruption, if Blaster didn't infect your network, it probably didn't affect you. Nevertheless, we kept our eyes on the indicators for signs of Internet distress in the hours and days that followed. After several days, trouble finally appeared, but Blaster didn't cause it.

Internet distress

At 4:11 PM EDT, Internet watchers saw signs of significant distress. The normal BGP chatter jumped several notches as border routers across the globe relayed the news of unreachable networks. Figure 1 shows the increase in route withdrawals from several sources beginning just after 20:00 Greenwich Mean Time (GMT). In a matter of minutes, more than 1 percent of the Internet was unreachable. As news of the blackout spread, we saw that the observed distress was not the result of an attack on the Internet itself, but one of the many secondary effects of the power-grid failure.

As the blackout continued, routing tables shrank as hundreds and then thousands of networks went offline. Figure 2 shows that several

thousand networks were withdrawn from global routing tables at the network reachability low point, which lasted until about midnight. From that point, networks slowly came back online over the next 24 hours, and the routing tables grew back to their normal levels. As you might expect, unreachable networks were concentrated in parts of Canada and the Northeastern US, particularly the New York City metropolitan area and Toronto. We observed a similar decline and rebound in network reachability from the Slammer worm, when routing table sizes dropped by about 3,000 networks and returned to normal levels after about 19 hours.[2]

Internet resilience

Packet-probe measurements during the blackout indicated that among the networks that remained online, there was no significant increase in end-to-end latency. This demonstrates resilience at the Internet's core, even when its edges were in distress. The lack of any major backbone failures prompted several news sources to report that the Internet was unaffected by the power failure. However, when thousands of networks disappeared for a 24-hour period, we lost services and business transactions that we could not measure with packet probes.

The availability of multiple metrics that don't always agree shows that measuring Internet performance is a multidimensional challenge and points to a need for several standardized measurements that would let us accurately gauge Internet health. Just as weather stations measure temperature, pressure, and humidity, and medical doctors check vital signs, including blood pressure, pulse, and respiration rate to monitor patient health, we need a set of standardized metrics that we can measure from multiple vantage points.

Based on the blackout experience, measurements from BGP and packet probes are a good start, but we also need other methods that let us accurately gauge the effect of traumatic events on Internet services and traffic. Such methods would not require special instrumentation because we could derive useful metrics from existing audit tools such as Web server logs, router flow logs, and intrusion detection systems. We could use any network device or service that logs its own activity as a local Internet sensor. For a global view, we could fuse local measurements into metrics that dampen isolated local effects and strongly indicate widespread phenomena.

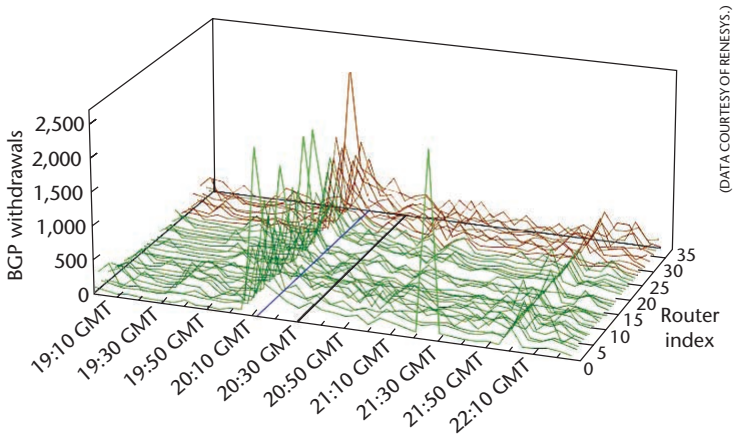

Figure 1. Border Gateway Protocol (BGP) route withdrawal spike. Border routers from all over the world saw the increase in withdrawals within a minute of the power outage. Spikes on a few border routers typically reflect local routing problems, but strongly correlated spikes in BGP activity are a sign of global Internet distress.

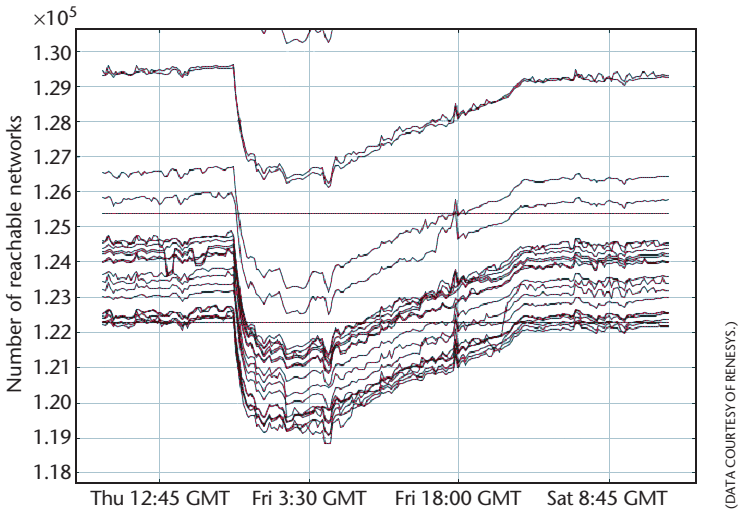

Figure 2. Border Gateway Protocol (BGP) table sizes reflect the loss of thousands of networks. More than 3,000 networks were unreachable for more than 8 hours following the blackout. Gradually, routing tables returned to normal size on Friday, 15 August as power slowly was restored.

Whether the 2003 blackout was just a glimpse of things to come as infrastructures become more interdependent or a once-in-a-generation event, we must gather as much information as we can when failures do occur. If there is a silver lining to this cloud, it is that with each blackout, worm epidemic, or other infrastructure catastrophe, we have an opportunity to discover more about our complex systems and improve their survivability.

Securing our critical infrastructures will require better understanding, and better understanding comes only

from data that we can study from an historical context. At the Institute for Security Technology studies, we plan to start by defining BGP-based routing metrics, including a global routing instability index and a global reachability index. These metrics will let us compare events such as the Northeast blackout, 9/11, the Code Red worms, and the Slammer worm to quantify their effects and use them as baselines for future events. □

References

1. J. Cowie et al., "Global Routing Instabilities During Code Red II and Nimda Worm Propagation," 2001; www. Renesys.com/projects/bgp_instability.
2. T. Griffin and M. Mao, "Interdomain Routing Streams," *Proc. Workshop on Management and Processing of Data Streams* (MPDS 03), AT&T Labs-Research, 2003; www.research.att.com/conf/mpds2003/schedule/griffinM.pdf.

Dennis McGrath *is a senior research engineer at the Institute for Security Technology Studies (ISTS) and the Thayer School of Engineering at Dartmouth College. His research interests include interdomain routing measurement, Internet health data correlation, and real-time simulation of cyber attacks. He earned his BS and MA from Rutgers University. Contact him at dennis.mcgrath@dartmouth.edu.*

How do we measure it?

How do we quantify the effects of traumatic events on the Internet? The most popular Internet performance measurement method is active measurement using probe packets. Probing using `ping`, `traceroute`, or similar "echo" applications can provide packet loss and latency data between two points. Several Web sites publish Internet performance measurements derived from systematic packet probes along fixed routes, but these are localized measurements, and it is difficult to infer global trends from local probes.

For a more comprehensive, global view of the Internet, we can derive performance metrics by listening to "conversations" between routers, specifically, analysis of border gateway protocol (BGP) messages. BGP is the exterior gateway protocol that makes inter-domain routing possible. Taken from multiple vantage points, these messages provide insight into global Internet activity.

Every domain or autonomous system (AS) connected to the Internet learns information about other networks by exchanging route information with neighbors. An AS might have a single neighbor, such as an upstream ISP, or hundreds of neighbors in the case of tier1 transit (backbone) networks, such as AT&T, Sprint, and Global Crossing. BGP lets AS neighbors share information about network reachability.

BGP conversations between border routers consist of two kinds of route messages: announcements and withdrawals. An announcement is a network reachability message, which specifies networks (identified as blocks of IP addresses) that are reachable via a particular route. Very often, a router changes its preferred route to a particular network and issues a new announcement. A withdrawal is a notification that a network is no longer reachable by any route via that neighbor.

Border routers build a BGP routing table, or route information base (RIB), from accumulated BGP announcements and withdrawals from all neighbors. Depending on the number of neighbors it has, a border router could have several routing options for each destination network. From its BGP table, the border router chooses its preferred route based either on a selection algorithm or the routing policy specified by the router administrator. The number of networks in the table increases as the Internet grows and, today, a border router with a global routing table has about 120,000 to 130,000 networks in its RIB.