Impact of the 2003 Blackouts on Internet Communications

Preliminary Report

James H. Cowie, Andy T. Ogielski, BJ Premore, Eric A. Smith and Todd Underwood

Renesys Corporation
www.renesys.com

November 21, 2003
(updated March 1, 2004)

Abstract

In August 2003, electric power outages affected 50 million people in the Northeastern US and Canada, causing economic losses estimated to exceed $5 billion. The outage was not only the largest in US history, but also the first of its scale since the rise of the commercial Internet and the World Wide Web.

In this report we examine the impact the blackout had on Internet connectivity and traffic routing in the region, as recorded in real time in data collected from Internet routers worldwide. We find that Internet connectivity in the blacked-out region was far more seriously affected than has been publicly revealed.

While the very largest provider networks (the Internet backbones) were apparently unaffected by the blackout [KN1, KN2], many thousands of significant networks and millions of individual Internet users were offline for hours or days. Banks, investment funds, business services, manufacturers, hospitals, educational institutions, internet service providers, and federal and state government units were among the affected organizations.

The scale and duration of the Internet connectivity outages recorded during the blackout strongly suggest that without additional investment in higher-quality interconnection and backup power at its edges, the Internet will be in no shape to supersede the telephone network as the nation's primary communications infrastructure.

This preliminary report is organized as follows: The main section is a summary of findings written with a minimum of technical language. Two appendices provide technical information concerning the methodology and a comparative analysis of Internet connectivity during a country-wide blackout in Italy on September 28, 2003 that reveals some interesting qualitative similarities despite significant differences in Italian Internet architecture.

Summary of findings

This report presents an analysis of the reliability of Internet connectivity to customer networks during the wide-area failure of the electric power grid in Northeastern US and Canada on 14–16 August, 2003.

According to our estimates, on August 14, 2003 there were over 9,700 globally advertised customer networks in the geographic area affected by the August blackout. According to the ARIN Internet Routing Registry [ARIN], they belonged to over 3,500 business entities and other private and public organizations, whose size and/or significant dependence on Internet communications warranted that routes to their own networks were advertised throughout the global Internet. This report shows that during the blackout the fate of these networks was as follows:

- 3,175 networks suffered from abnormal connectivity outages. Of those, more than 2,000 networks suffered severe connectivity outages for longer than 4 hours, and over 1,400 networks for longer than 12 hours (some even longer than 48 hours).
- The networks suffering from abnormal connectivity outages belonged to over 1,700 organizations (business entities, government, education and other institutions). More than 1,000 organizations had outages of *all* of their networks lasting longer than 4 hours.
- Nearly 50% of Autonomous Systems (organizational entities involved in global Internet routing, typically corporations) lost connectivity to some or all of their networks in the blackout area.

Background

At 16:10 EDT on August 14, a cascading power grid failure rippled across a large sector of the Northeastern US and Ontario, Canada. Within two minutes, the outage was complete across the region. In contrast, restoration was a slow, incremental process. Power was completely restored in New York at 21:00 EDT on Friday, August 15, and in the entire blackout area by Sunday, August 17.

In order to determine the impact of power outages on Internet customers, it is sufficient to analyze Internet routing to customer networks. If there were no routes from any location worldwide to a given network in the blackout area, then that network was cut off from the Internet, and any packets sent to it were lost. Renesys Corporation continually collects second-by-second Internet routing

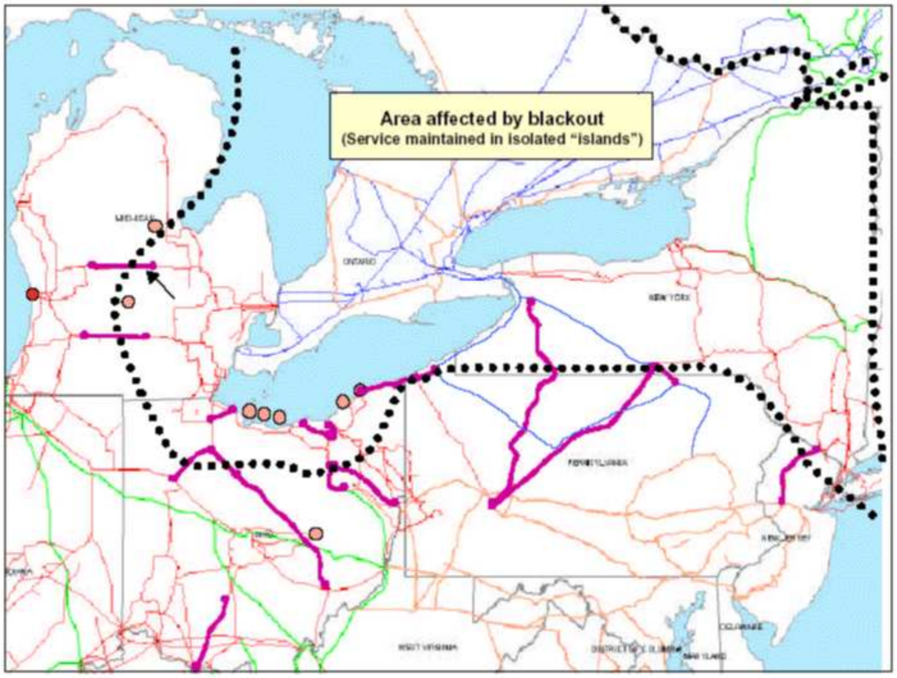

Figure 1: 14 August 16:13 EDT—cascading power failure complete. Blacked out areas are within the dotted line. Source: [TF1,TF2].

data from dozens of locations worldwide, and correlates the data to provide both real-time and forensic analysis.

A network is represented by a *network address*, which is an aggregated range of IP addresses. A *globally advertised network* is a network whose address is made visible by its owner to the world, in the sense that under normal conditions most border routers in the world would have a path for forwarding IP packets to this network in their routing tables. Larger enterprises connected to the Internet usually own a number of these globally advertised networks.

Customers connect to the global Internet through one or more Network Service Providers (NSPs). In the simplest case, a single NSP takes responsibility for making the customer's networks visible to the rest of the Internet by advertising them exclusively on the customer's behalf. For increased redundancy, and to retain more control over how their networks are advertised, a customer may contract for Internet service from more than one NSP.

These *multihomed* customers are themselves a direct part of the global routing infrastructure; they, and not their NSPs, send the original advertisements for their networks out into the Internet. To achieve this, a customer organizes its networks into one or more Autonomous Systems (ASes) that are registered with a regional Internet Routing Registry. NSP backbones are Autonomous Systems as well. An AS is thus an autonomous component of the global routing infrastructure, with *border routers* on its perimeter. The border routers exchange routing information messages with neighboring ASes using the Border Gateway Protocol (BGP). BGP is the common mechanism that connects all Autonomous Systems together into one global Internet.

Because routing information messages propagate throughout the Internet in this way, it is possible to learn about network reachability by collecting routing messages from virtually any border router in the Internet. In particular, analysis of messages arriving at multiple, geographically diverse border routers can detect temporally correlated "waves" of withdrawal messages, which indicate the failure and disconnection of a customer network from the Internet. This type of analysis is the approach taken in this report.

For brevity, we say that a customer network has an *outage* when it is disconnected from the Internet. Note that a network outage during an electric power failure may be due to a lack of reliable backup power on the customer side, on the NSP side, or somewhere along the physical links connecting them; therefore network outage data alone cannot be used to assign blame conclusively to any of these failure points.

While we exercised considerable caution in data analysis and validation, it is important to understand that some inaccuracies may have existed in data sources over which we had little or no control. In particular, there may be errors due to incorrect identification of network owners in Internet Routing Registries, imprecise geolocation of networks, or lack of high geographic resolution of the boundaries of blacked-out regions (there were pockets where power was available). We sampled the list of networks for additional verification, and we believe that with data available at the time of writing the report is reasonably accurate.

Network Outage Statistics

We generated a list of all networks *inside the blackout area*, including their origin AS number, owner organization name and address and ZIP code according to the ARIN database [ARIN]. The list includes both customer networks connected directly via their NSPs, and customer networks organized into Autonomous Systems. For each network on the list, we computed all its outages intersecting the time window from just before the cascading blackout (Thursday 14 August 16:00 EDT) to Monday 18 August 16:00 EDT. This time window was chosen because dozens of customer networks had outages lasting several days from the onset of blackout. Short transient outages that can occur under normal conditions were discarded, and only outages longer than 5 minutes were considered.

In order to obtain a quantitative characterization of the impact of blackouts on Internet customers, we use the following metrics:

- Number and durations of outages of customer networks located in the blackout area.
- Numbers of Autonomous Systems located in the blackout area, broken down into categories according to severity of their network outages, and according to patterns of their connectivity to the Internet.
- Numbers of business entities and other institutions located in the blackout area, broken down into categories according to severity of their network outages.

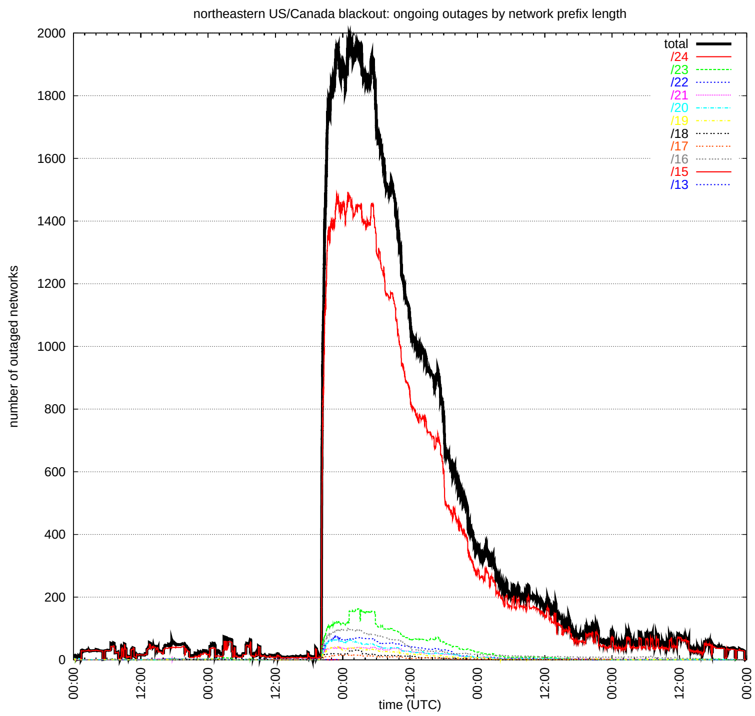

Figure 2: The number of *concurrently ongoing* network outages in the blackout area are shown at every moment between the beginning of August 13 (00:00 UTC) to the end of August 17 (24:00 UTC). The black line represents the total number of outages. The colored lines represent the contribution of outages of networks grouped according to their size (number of IP addresses in a network).

The distributions of certain metrics during the blackout turn out to have very characteristic shapes, present in the Italian blackout (September 28, 2003; see Appendix B) as well as in the US/Canadian event.

Figure 2 shows the time course of all ongoing network outages during the blackout. As elsewhere in this report, only networks located in the blackout area are considered. The interpretation of Figure 2 is as follows: for any date and time value on the horizontal x-axis, the corresponding height of the plot is equal to the number of concurrently ongoing network outages. As time progresses, some new outages may begin while others end. The extremely steep rise of the plot begins on August 14, 20:11 UTC (16:11 EDT), exactly the time of cascading power grid failure.

The colored plots in Figure 2 represent contributions of outages of networks grouped according to their size. A network's size is the number of IP addresses it contains, and in technical notation "/24" identifies networks with capacity of 256 IP addresses, "/23" means 512 addresses, and so on, doubling with each step. Notice that the smallest networks (/24s) contribute the vast majority of outages. However, in the US it is for historical reasons that /24s are the most common, and their small size does not necessarily indicate that the network belongs to a small organization.

The *baseline noise* in the plot before the beginning of the blackout shows ongoing network outages due to other causes. MS Blaster worm was active in that period, and may have contributed to higher

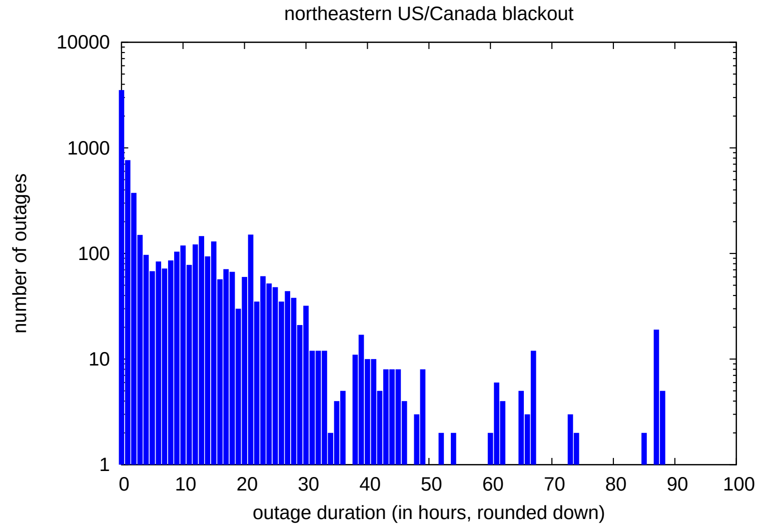

Figure 3: Distribution of the number of network outages versus outage duration (rounded down to full hours) during the blackout.

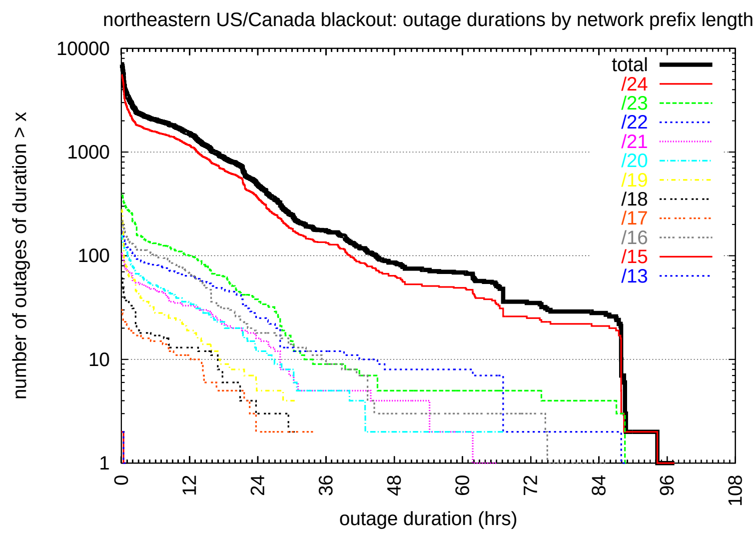

Figure 4: Cumulative distribution of the number of network outages. On the horizontal axis is outage duration T, on the vertical axis is the number of recorded outages of duration longer than T.

level of network outages due to maintenance and other collateral damage.

While Figure 2 shows the number of ongoing outages at any time, it does not indicate outage durations. They are presented in Figures 3 and 4. Only abnormally long network outages are shown here; as elsewhere in this report, outages shorter than 5 minutes have been discarded.

Network outage durations shown in Figure 3 are rounded down to full hours (so, for example, an outage lasting between 2 hours and 3 hours is rounded to 2 hours). The distribution of outage

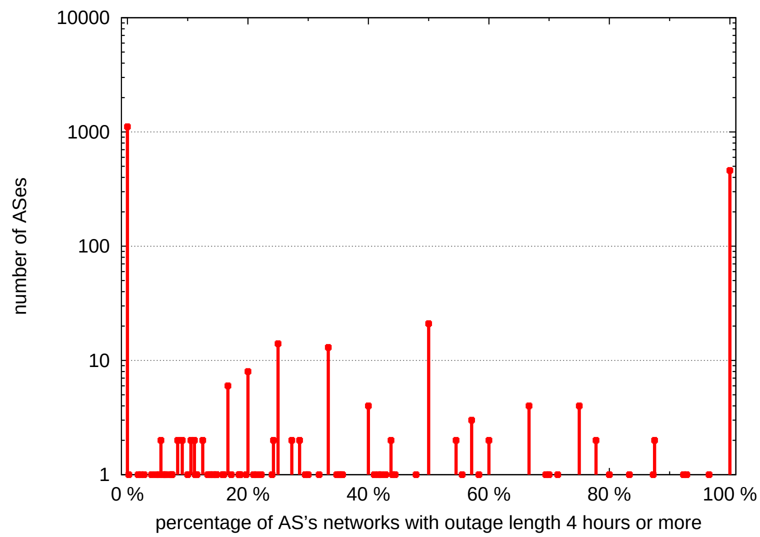

Figure 5: Bar chart showing the number of Autonomous Systems versus the percentage of their networks that had outages longer than 4 hours (outages shorter than 4 hours are ignored). The peak at 0% has 1,110 ASes, the peak at 100% has 460 ASes, and there are only 163 ASes in the middle between the two peaks.

durations has a characteristic shape: A high peak of relatively shorter outages, then a rounded bump continuing into a long sparse tail. 2,270 outages in the bump and tail were longer than 4 hours, with most in the range from 4 to 24 hours.

In order to show the number of outages *longer* than a given number of hours, Figure 4 presents the cumulative distribution of network outages.

It is worth noting that, except for the magnitudes, the shapes of the distributions in Figures 3 and 4 are similar to their counterparts obtained from the Italian blackout data (see Appendix B).

Network Outage Distributions in Organizations

Next, we turn to the question of how the impact of network outages was distributed among the Autonomous Systems that owned networks located in the blackout area. There is no unique methodology to address this question, as it can be approached from many directions. A reasonable measure of the severity of outages in a given AS is the percentage of networks in the AS that had outages longer than 4 hours. (We also investigated analogous severity measures for threshold times both shorter and longer than 4 hours, and the results remain qualitatively similar.)

Note that due to the high variability of outage durations, a simple (unconditional) average network outage time per AS is not a good severity measure when long-duration outages are of particular interest.

The distribution of the number of ASes according to the percentage of their networks that had outages longer than 4 hours has a striking shape: There are two large peaks at exactly 0% and 100%

corresponding to ASes that had *no* outages longer than 4 hours, and to ASes whose *all* networks had outages longer than 4 hours, respectively. In the middle between these two peaks are ASes in which some networks had outages, and some did not. Notice that the middle part of the distribution is not uniformly populated: There are more ASes with fewer than 50% of network outages.

Detailed examination of the list of networks and their origin ASes shown in Figure 5 reveals interesting trends concerning the nature of the organizations owning these networks. The number of networks in the ASes in the peak at 0% is 4,735, the number of networks in the peak at 100% is 997, and the number of networks in the middle between the peaks is 3,991. One immediately notices that the ASes with no outages (0%) tend to have *more networks per AS* than the ASes that lost every network (100%). A review of the datasets confirms that the ASes with no network outages tend to be larger organizations owning more networks, while ASes that lost every network (100%) tend to be smaller organizations with few networks. There are exceptions to this statement, but the trend is clear.

Observe that the Autonomous Systems in the middle of the distribution have the most networks per AS (3,991 networks and 163 ASes). Examination of these ASes reveals that majority of them are wide-area and regional Network Service Providers, and the networks originated in these ASes belong to their customers (business entities and various public and private institutions). The Network Service Providers with over 50 advertised networks located in the blackout area are listed in Table 1, together with the percentage of their networks suffering from outages longer than 4 hours (shorter outages are ignored), and average network outage duration (again, including only outages longer than 4 hours).[1] Dial-up, cable and other Internet access providers are not NSPs and are not included.

[1]Updated March 1, 2004 to filter scheduled maintenance outages reported to the authors.

AS Number	Provider AS	est. # networks in blackout area	% network outages $\geq$ 4 hr	average outage duration (hours)	comments
209	Qwest	109	6%	29	
237	Merit	94	35%	24	Michigan
577	Bell Adv. Comm	160	11%	17	Ontario
701	UUNET	266	41%	22	
813	UUNET	50	40%	13	Ontario
852	Telus	64	14%	16	Canada
1239	Sprint	61	29%	22	
1785	FASTNET	97	18%	12	NY, PA
2386	ATT Data Comm	136	7%	20	
3602	Sprint Canada	66	20%	30	
4471	MICA	55	87%	22	Michigan Internet
5650	Electric Lightwave	57	2%	13	Rochester NY
7018	AT&T Worldnet	328	35%	15	
7046	UUNET	66	24%	33	
12021	Hamilton Hydro	55	27%	16	Fibrewired, ON
14456	Peterborough Util	52	0%	—	Ontario
15290	Allstream	115	15%	18	was AT&T Canada

Table 1: Network Service Providers with 50 or more customer networks in the blackout area.

The largest percentages of network outages (75% to 96%) were suffered by customers of 8 NSPs, all but one (AS 4471) with fewer than 50 advertised networks in the blackout area, and in one university.

Next, we turn to the question of how the impact of network outages was distributed among the Internet customers that owned networks in the blackout area. This is a considerably more difficult question than the analysis at the resolution of Autonomous Systems described above. Using the data available to us, we cannot directly measure outage impact in terms of loss of productivity or revenue. However, we can compute a similar severity measure for each Internet customer (business entities, government, education and other institutions): the percentage of their networks that had outages.

We distinguish the network owners by their names and addresses as provided by the ARIN Internet Routing Registry [ARIN]. We notice that this introduces ambiguities when the networks actually

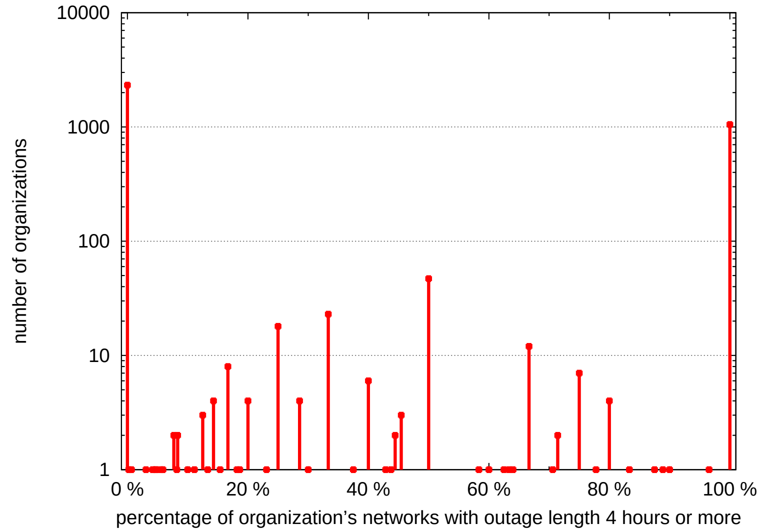

Figure 6: Bar chart showing the number of organizations in the blackout area (business entities, government, and other public and private institutions) versus the percentage of their networks that had outages longer than 4 hours.

belonging to the same organization are registered under different names and/or addresses (due to corporation name changes, acquisitions, entry errors, etc.) and vice versa, but short of an exhaustive manual verification it is difficult to correct the ARIN database errors, and therefore in the following no such corrections were made.

The distribution of the number of organizations according to the percentage of networks that had outages longer than 4 hours is shown in Figure 6, and is qualitatively similar to Figure 5 (for Autonomous Systems).

In Figure 6, the number of organizations in the peak at 0% is 2,321, the number of organizations in the peak at 100% is 1,051, and the number of organizations in the middle between the peaks is 185. The corresponding numbers of networks owned by these three groups are 6,135 (0%), 1,566 (100%) and 2,022 (the middle).

The trend distinguishing the organizations in these three groups is similar to that observed in the corresponding distribution for ASes: On the average, the organizations that had 100% network outages have the fewest networks per organization (1.5), the organizations that had 0% of network outages had 2.6 networks/organization, while the organizations in the middle had the largest average number of networks per organization, 10.9.

Examination of the list of *organization names* in each of the three categories, however, did not easily reveal any clear trends. More research is required to determine what correlations may exist between the type of business, its revenue, the importance of Internet for conduct of its operations, the number of business locations, etc., and the severity of network outages during the blackout.

The bar charts shown in Figures 5 and 6 have a peculiar shape, with large peaks at 0% and at 100%,

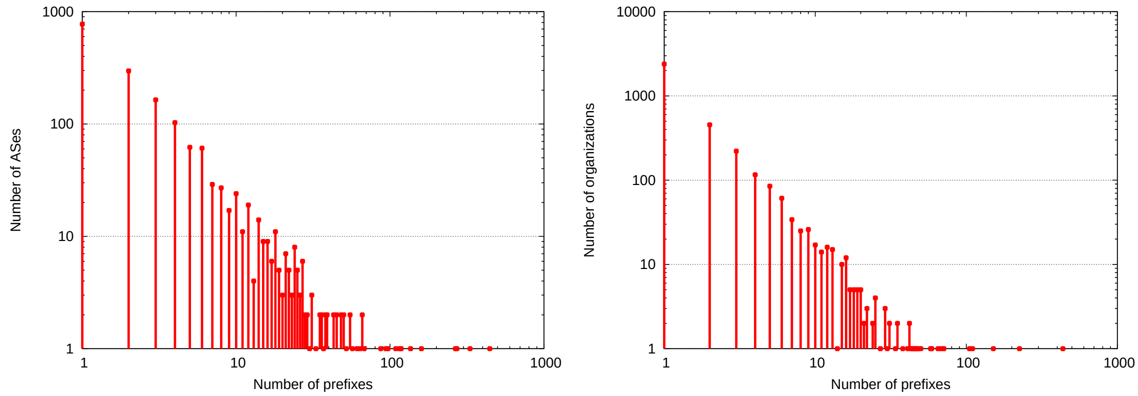

Figure 7: Left: Distribution of the number of Autonomous Systems versus the number of their originated networks. Right: Distribution of the number of organizations versus the number of their networks. Only networks localized within the US/Canada blackout area are counted.

and a small bump in the middle between the peaks. It turns out that the reason for this is a *very high variability* of AS and organization sizes, as measured by the number of networks they own: the vast majority are small, a few are extremely large, with not many between the extrema (mathematically speaking, the distribution of sizes can be approximated by a power law). This is shown in Figure 7.

The qualitative explanation is based on the observation that for the distribution of AS (or organization) sizes shown in Figure 7, if one assumes an approximately constant probability of network outage, the resulting distribution of the number of ASes (organizations) according to the percentage of their network outages has a qualitatively similar shape to that in Figures 5 and 6.

The preceding analysis shows that organizations that advertise larger numbers of networks tended to be more resilient to power outages.

Edge Customer ASes

Finally, we investigate the trends among edge customer ASes, that is, customer ASes that do not transit externally originated traffic. As explained earlier, organizations that connect to the Internet in this way maintain control over how their networks are advertised to the global Internet, and may have multiple Network Service Providers.

The data, shown in Figure 8, indicate that networks owned by multihomed customers (with two or more NSPs) generally tended to be more resilient to failure than those of single-homed customers. The obvious outlier, seen in the right panel of the figure as having 9 upstreams, corresponds to networks owned by a single midsize Internet Service Provider serving businesses in Manhattan.

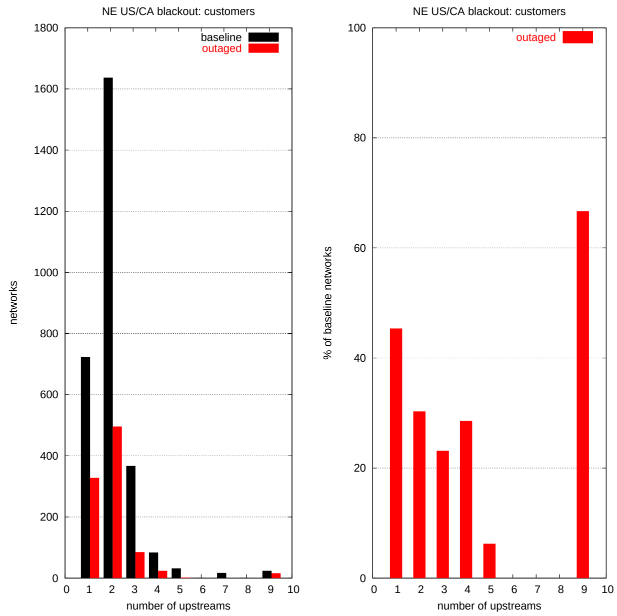

Figure 8: Breakdown of the population of networks originated by edge customer ASes by the number of their upstream network service providers. Only the networks in the blackout area are included. Left panel: Comparison of the total number of globally advertised networks prior to the blackout (in black) to the number of networks that suffered from outages (red). Right panel: Numbers of networks that had outages shown as a percentage of the total prior to the blackout.

Conclusions

Reports that the Internet was generally unaffected by the August blackout [KN1,KN2] significantly understated the event's impact on the nation's emergent critical communications infrastructure.

The very largest provider networks (the Internet backbones) were reportedly unaffected by the blackout, but thousands of corporate and other institutional networks and millions of individual Internet users were offline for hours or days. Banks, investment funds, business services, manufacturers, hospitals, internet service providers, and federal and state government units were among the affected organizations.

The total number of US businesses and other institutions affected by the power outages may never be known with certainty. The methods used in this report tend to significantly undercount blackout-related Internet outages, because we examined only those networks that are individually advertised world-wide. Small businesses that rely on dial-up or DSL connectivity, or whose networks are lumped together with their providers' networks before they are globally advertised, probably had significant power and/or network outages of their own that would not appear in the global routing datasets that we examined.

On the brighter side, the specific effects of the blackout on Internet availability were geographically well-localized, and we found no evidence of cascading failures affecting global Internet stability.

In general, organizations that advertise larger numbers of networks tended to fare better, as did better-connected organizations who could afford to purchase service from a diverse set of providers.

One feasible interpretation of our data is that Internet customers that own and operate well-connected networks also tend to have the resources to afford higher quality backup power, as well as skilled personnel to restore network services faster. In other words, the economics of end-user networks, rather than core networks, were the primary determinant of Internet survivability during the blackout.

The scale and duration of the outages we measured during the blackout strongly suggest that without additional investment in higher-quality interconnection and power at its edges, the Internet will be in no shape to supersede the telephone network as the nation's primary communications infrastructure.

About the authors

Renesys Corporation, with main offices in Hanover, NH, delivers the first commercial global Internet monitoring services to enterprises and government customers. Company information, whitepapers and research reports are available on the web, http://www.renesys.com.

Appendix A. Methodology

US/Canada blackout: Sequence of events and affected areas

The 14–16 August blackout was the largest in US history, affecting an area inhabited by approximately 50 million people. According to [TDW], in comparison the 1965 blackout affected 25 million people, the 1977 blackout—9 million people, and the 1996 West Coast blackout affected 4 million people.

A sequence of power grid events provides the context for the Internet outages. The times of blackout onset events are extracted from [TF1,TF2]. In contrast to the extremely rapid progress of cascading power failures, wide-area power restoration was a very complex, long-lasting, incremental process. A partial sequence of power restoration compiled from reports issued by regional transmission operators [IMO, ISONE, MISO, NYISO, PJM] and the press is shown in Table 2.

15:05 – 16:10,	Thursday, Aug 14	Multiple transmission line disconnects in Ohio.
16:10 – 16:13,	Thursday, Aug 14	Cascading blackout.
05:30,	Friday, Aug 15	56 percent of the NY State load restored.
12:00,	Friday, Aug 15	Electric service restored to all PJM customers (PA). Ontario electricity generation restored to over 50 percent capacity.
16:00,	Friday, Aug 15	Over 60% of power restored to areas of central and northern Michigan, all of northern Ohio and approximately one-third of the Detroit region.
21:00,	Friday, Aug 15	New York City power restored.
22:45,	Friday, Aug 15	NY State power restored in remaining pockets of affected areas.
24:00,	Friday, Aug 15	Power restored to nearly all customers in New England.
early,	Saturday, Aug 16	Transmission system in Southwest Connecticut restored.
	Sunday, Aug 17	Power restoration in the Midwest nearly completed.

Table 2: Abbreviated timeline of the US/Canada blackout. All times are EDT = UTC – 4:00.

Data sources and methodology

The raw input dataset consists of all BGP routing messages received from approximately 30 BGP routers. These default-free (full-table) BGP sessions have been maintained for many months with remote routers in diverse geographic locations and diverse ASes in Asia, South Africa, Europe and North America. All routing messages are timestamped and archived in Renesys databases. They provide second by second, multiple views on the routes to all networks in the world, and in particular to the networks in the blackout areas.

Geolocation of networks was initialized with netblock owner addresses and ZIP codes provided by the ARIN and RIPE Internet Routing Registries [ARIN,RIPE]. Translation of ZIP codes to latitude and longitude was performed with a Renesys software package. It is known that the Routing Registries contain errors, for instance networks actually belonging to the same organization are occasionally registered under different names and/or addresses, corporation name changes, mergers and acquisitions are not always accounted for, there are occasional typographic errors, etc. Short of an exhaustive manual verification it is difficult to correct the ARIN database errors, and therefore no such corrections were made except in a few obvious cases.

In the first step, two datasets restricted only to networks located in the geographic blackout areas (US/Canada and Italy) were extracted from the BGP routing tables, at a time approximately an hour before the cascading power failures. Given that some BGP routers announce slightly different sets of networks, these datasets were further restricted to contain only the networks advertised by at least five distinct routers. A small number of anomalous routes to networks smaller than /24 (in the /25–/32 range) were rejected.

The Italian network dataset thus contains all networks located in Italy that were present in the global routing tables of five or more Autonomous Systems just prior to the September 28 blackout.

In order to construct the US/Canada network dataset, in addition we had to estimate the boundary of the geographic blackout area. At the time of writing of this report we had no high resolution (county level or better) borders of the blacked-out areas. Therefore, we took the coarse-resolution boundary of the blackout area from the US/Canada Task Force Report (see Figure 1), then enclosed the curvilinear blackout boundary in Michigan, Ohio, Pennsylvania and northernmost New Jersey in a bounding box in spherical coordinates, and extracted all networks contained in this bounding box. Subsequently, we merged this set with the set of all networks located in New York State and in the Province of Ontario, Canada. Therefore, the final US/Canada network dataset contains all networks within the area outlined in the Task Force Report, and in addition some networks in eastern Michigan, central Ohio, and northern parts of Pennsylvania and New Jersey. The resulting roundoff error can be corrected in the future once detailed blackout boundary coordinates become available. Note that inclusion of additional networks from outside of the probable blackout area tends to reduce the percentage of outages caused by the blackout, in agreement with our intention not to exaggerate the impact of the blackout on Internet connectivity.

In the next stage we generated datasets of all outages for the networks contained in the restricted US/Canada and Italian datasets, in the time windows beginning before the blackouts and ending after complete power restoration (see below). A network outage event is considered to begin when at least $P = 5$ Renesys peer routers signal withdrawals for the network address (also known as a network *prefix*) such that no two of these withdrawals are separated by more than 30 seconds. In addition, no new advertisements may be seen during the entire “outage start” interval.

The outage is considered to have started when the P-th withdrawal arrives, and continues while at least P peers continue to see the prefix as withdrawn (these peers need not be the same as the original P peers that began the outage), ending as soon as fewer than P peers see the prefix as withdrawn.

Only outages longer than 5 minutes are counted. Based on our BGP monitoring experience, such a definition exploits the high correlation of withdrawal times by diverse peers, suppresses transient withdrawals that commonly accompany route changes, and in practice assures that the network was withdrawn by its origin AS or by all of its upstream ASes.

It is noted that since the recommended practice is not to redistribute IGP routes into eBGP, a BGP border router that remains on backup power may not withdraw statically configured routes to internal or customer networks that did actually lose power. This is another source of undercounting the network outages due to the blackouts.

In this report we do not assign significance to situations when a route to a highly aggregated network prefix was not withdrawn, while more specific networks were withdrawn. In theory, when more specific routes are withdrawn, but a route to the containing /8 or /13 is not, the more specific remains technically reachable. However, in practice the large aggregates such as 12.0.0.0/8 (AT&T Worldnet) or 35.0.0.0/8 (Merit) are virtually never withdrawn when more-specifics are; moreover, multi-level aggregation hierarchies are not very common in the blackout datasets.

The estimated numbers of networks and Autonomous Systems physically located in the blackout areas are shown in Table 3. Physical networks of certain ASes may extend beyond the blackout area, such as AS 701 and other wide-area NSP backbones.

	US/Canada	**Italy**
# ASes originating networks in the blackout area	1733	234
# networks in the blackout area	9723	686
# edge customer ASes in the blackout area	975	123
# networks in edge customer ASes in the blackout areas	2899	162

Table 3: Sizes of network and AS datasets used in this report.

Appendix B. Italian blackout

Italy, with a population of 57M people, experienced what appears to be the largest European blackout since WWII on Sunday, September 28, 2003. The entire country was blacked out after the cascading failure, with the exception of the island of Sardinia. The blackout began at 01:25 UTC and spread throughout the country in about three minutes [UCTE]. Restoration began in northern regions around 05:00 and continued until it was completed at about 21:00.

There are some substantial differences in the statistics and architecture of the Internet in North America and Italy. They are interesting because the two blackouts affected comparable populations—approximately 50 million people each—in economically advanced countries. The striking differences between the two regions are:

- The number of globally advertised networks is over 14 times larger in the North American blackout area than in all of Italy.
- The number of Autonomous Systems that are either fully or partially inside the North American blackout area is over 7 times larger than the number of ASes in all of Italy.
- The degrees of route aggregation (distributions of routable network sizes) are qualitatively significantly different. North American networks are overwhelmingly dominated by /24s, while Italy has a more even distribution (/19s dominate, with /24s, /16s and /20s close behind). This reflects the much earlier (pre-CIDR) deployment of the Internet in the USA. Figure 9 illustrates this comparison. Notice that although Italy has more /19s than /24s, a higher fraction of /24s went out. This can also be seen in Figures 10 and 11.

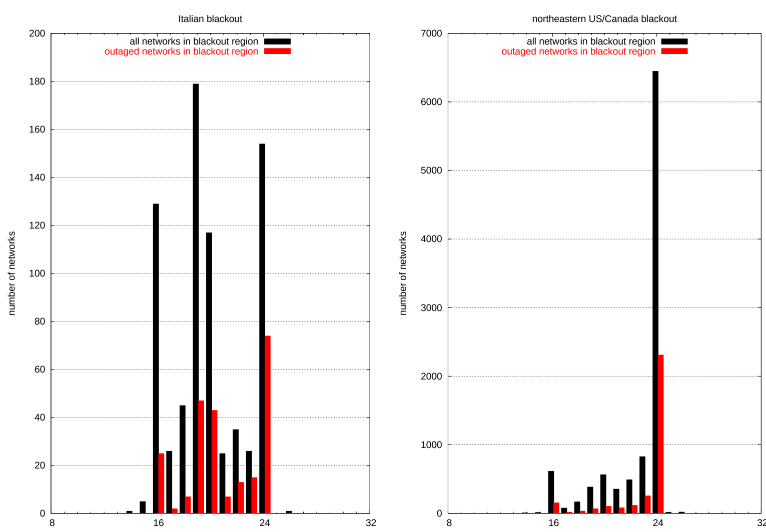

Figure 9: Left panel: Number of networks suffering an outage compared to the total number of networks in the Italian blackout area, broken by prefix length. Right panel: Analogous bar chart for the US/Canada blackout, note significantly larger number of networks, and a distinct distribution of networks by prefix length.

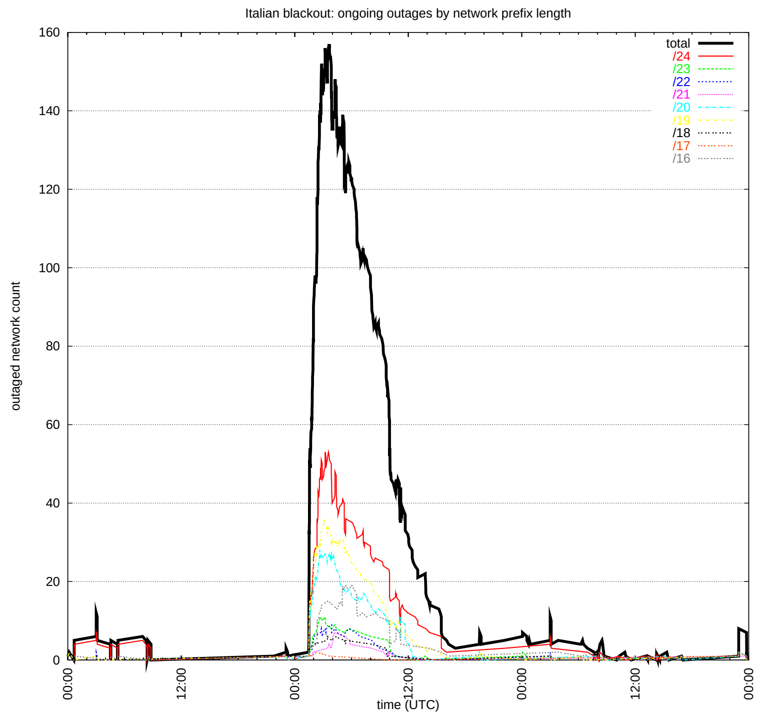

Figure 10: Total number of all concurrently ongoing outages of networks located in Italy before, during and after the blackout. Long duration outages dominate in this plot. The left end of the time axis is the beginning of September 27 (00:00 UTC).

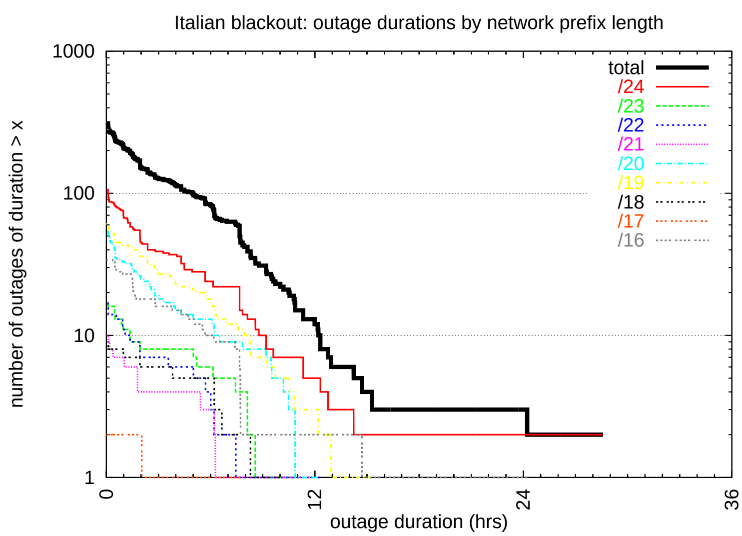

Figure 11: Cumulative distributions of Italian network outage durations by network prefi x length. On the x-axis is outage duration, on the y-axis the number of recorded outages of duration longer than x.

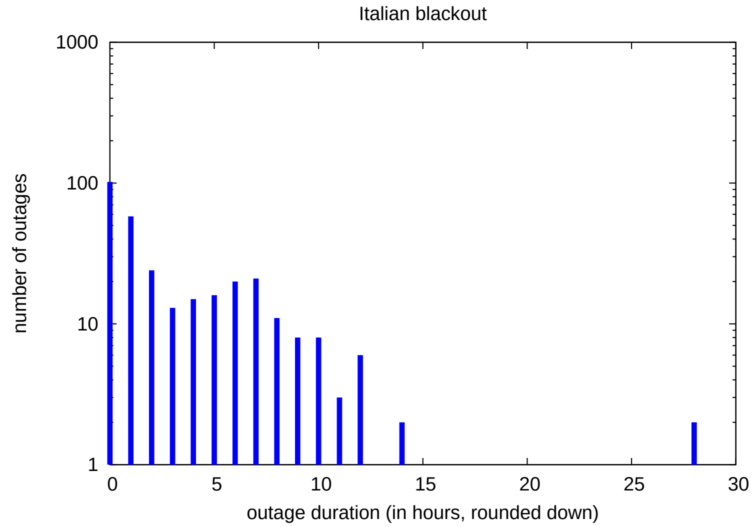

Figure 12: Distribution of durations of all Italian network outages.

However, despite these differences—most notably the vast disparity in wide-area network sizes in the two blackout areas—there were similarities in the statistics of network outages. Specifically, compare the shape of Figure 12 with that of its US/Canada counterpart (Figure 3). These similarities may indicate the existence of more general statistical laws describing the expected number of connectivity failures and their time courses under wide-area network failures due to blackouts or other disasters. This is a suggested subject for future research.

The analysis of the blackout is performed using the same methodology that has been applied to the analysis of the US/Canada blackout. Therefore, it is summarized with a minimum of comments for the sake of brevity.

At the time of writing, the available reports on the Italian power grid collapse were considerably less detailed that those for the US/Canada event. Approximate power restoration times were obtained from the Italian and international press.

Similarly as in the case of US/Canada blackout, when customer ASes are characterized by the number of their upstream network service providers, the networks in an AS with a larger number of upstreams tend to be more resilient to failure. The Italian data shows a monotonic trend towards higher resilience with the number of providers increasing from 1 to 3; however, there is an insufficient number of customer ASes with more providers than two to draw firm conclusions.

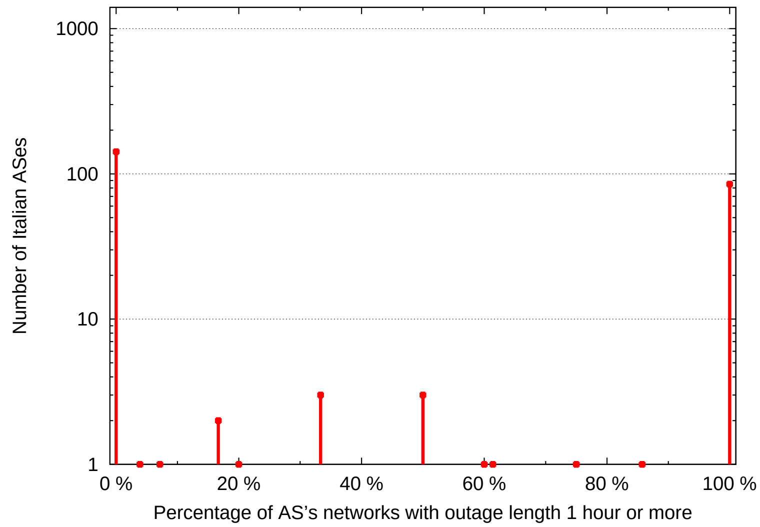

Figure 13: Bar chart showing the number of Autonomous Systems versus the percentage of their networks that had outages longer than 1 hour (outages shorter than 1 hour are ignored).

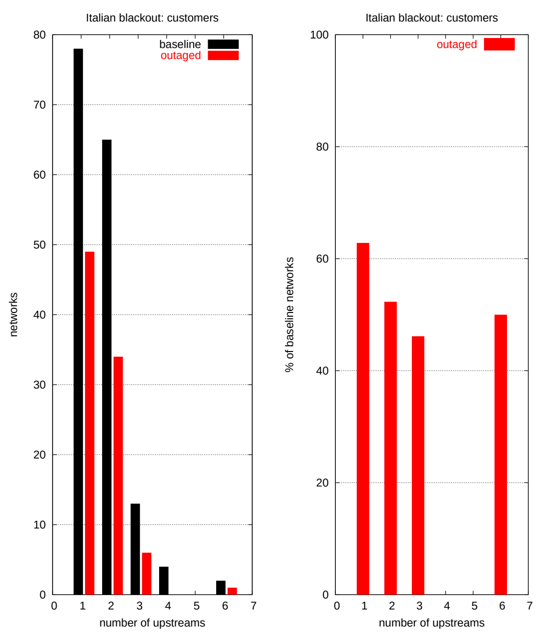

Figure 14: Breakdown of the population of networks originated by edge customer ASes by the number of their upstream network service providers. Only the networks in Italy are included. Left panel: Comparison of the total number of routable networks prior to the blackout (in black) to the number of networks that suffered from outages (red). Right panel: Numbers of networks that had outages shown as a percentage of the total prior to the blackout.

References

ARIN ARIN Internet Routing Registry.
http://www.arin.net/

IMO Preliminary IMO Board Report to Market Participants Following August 14, 2003 Market Suspension. Independent Electricity Market Operator (Canada).
http://www.theimo.com/imoweb/pubs/corp/reports/br_MS-2003aug14-20030908.pdf, see also http://www.theimo.com/imoweb/pubs/pressReleases/pr_powerRestorationMilestones.pdf

ISONE Independent System Operator New England News and Press Releases.
http://www.iso-ne.com/iso_news/newnews.html

KN1 Internet Performs Reliably During Northeast Power Outage According to Keynote. Keynote Press release, 14 September 2003.
http://www.keynote.com/news_events/releases_2003/03august14.html

KN2 Internet Performance Holds Steady So Far, According to Keynote. Keynote Press release, 18 September 2003.
http://www.keynote.com/news_events/releases_2003/03august18.html

MISO Midwest ISO News.
http://www.midwestiso.org/newsroom.shtml

NYISO August 14, 2003—System Disturbance. New York Independent System Operator.
http://www.nyiso.com/topics/outage2003_info/index.html

PJM PJM Interconnection regional transmission organization. News.
http://www.pjm.com/contributions/news-releases/2003/2003.html

RIPE RIPE Internet Routing Registry.
http://www.ripe.net/

TDW Transmission and Distribution World. September 1, 2003.
http://tdworld.com/ar/power_fifty_million_people/index.htm

TF1 August 14, 2003 Outage Sequence of Events. US/Canada Power Outage Task Force. 12 September 2003.
http://www.ferc.gov/cust-protect/moi/09-12-03-blackout-sum.pdf

TF2 U.S.-Canada Power System Outage Task Force Interim Report: Causes of the August 14th Blackout in the United States and Canada. November 19, 2003.
https://reports.energy.gov/

UCTE Interim Report of the Investigation Committee on the 28 September 2003 blackout in Italy. October 27, 2003.
http://www.ucte.org/news/e_default.asp