SANDIA REPORT
SAND2017-13442
Unlimited Release
July 2017

Modeling the Internet

Andjelka Kelic, Michael Mitchell, Donald Shirah

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: http://www.ntis.gov/search

SAND201X-XXXX
July 2017
Unlimited Release

Modeling the Internet

Andjelka Kelic, Michael Mitchell, Donald Shirah
Resilience and Regulatory Effects, Systems Research, Analysis, and Applications
Sandia National Laboratories
P. O. Box 5800
Albuquerque, New Mexico 87185-MS1137

Abstract

The National Infrastructure Simulations and Analysis Center (NISAC) has developed a nationwide model of the Internet to study the potential impact of the loss of physical facilities on the network and on other infrastructures that depend on the Internet for services. The model looks at the Internet from the perspective of Internet Service Providers (ISPs) and their connectivity and can be used to determine how the network connectivity could be modified to assist in mitigating an event. In addition the model could be used to explore how portions of the network could be made more resilient to disruptive events.

ACKNOWLEDGMENTS

Funding for this work was provided by the Department of Homeland Security (DHS) Office of Cyber and Infrastructure Analysis (OCIA). OCIA provides innovative analysis to support public and private-sector stakeholders' operational activities and effectiveness and to inform key decisions affecting the security and resilience of the Nation's critical infrastructure. All OCIA products are visible to authorized users at HSIN-CI and Intelink. For more information, contact OCIA@hq.dhs.govor visit http://www.dhs.gov/office-cyber-infrastructure-analysis.

TABLE OF CONTENTS

FIGURES

TABLES

EXECUTIVE SUMMARY

In support of the Department of Homeland Security (DHS), Office of Cyber and Infrastructure Analysis (OCIA), the National Infrastructure Simulation and Analysis Center (NISAC) at Sandia National Laboratories (SNL) is developing a national model of the Internet, Athena, to assess the performance of the nation's Internet infrastructure and how it may be impacted by natural disasters and other threats. The model looks at the Internet from the perspective of Internet Service Providers (ISPs) and their connectivity and can be used to determine how the network could be modified to mitigate the impacts of a disruptive event. It will also enable assessment of impacts on dependent infrastructure and determine areas of common vulnerability.

The initial version of the model includes the twelve largest US providers, plus two additional regional providers to better study the network interactions. Providers were selected based on rankings of direct and indirect customers from the Center for Applied Internet Data Analysis (CAIDA).

The network topology is based on embedding the Autonomous System (AS) business relationships into the topology. This topology allows for the use of graph theory and associated algorithms to determine whether or not the network is fully functional or has been disconnected at some point. A working network results in a single graph, while a disrupted network results in multiple disconnected graphs. The analyst can then explore the graph to determine where the network could be reconnected by changing a business relationship rule if possible, or adding additional connections to the network.

NOMENCLATURE

Abbreviation	Definition
AS	Autonomous System
ARIN	American Registry for Internet Numbers
BGP	Border Gateway Protocol
CAIDA	Center for Applied Internet Data Analysis
DHS	Department of Homeland Security
FCC	Federal Communications Commission
ISP	Internet Service Provider
NISAC	National Infrastructure Simulation and Analysis Center
OCIA	Office of Cyber and Infrastructure Analysis
SNL	Sandia National Laboratories

1. INTRODUCTION

In support of the Department of Homeland Security (DHS), Office of Cyber and Infrastructure Analysis (OCIA), the National Infrastructure Simulation and Analysis Center (NISAC) at Sandia National Laboratories (SNL) is developing a national model of the Internet, Athena, to assess the performance of the nation's Internet infrastructure and how it may be impacted by natural disasters and other threats. This capability will improve DHS's ability to assess Internet availability in the event of a disruption and to determine potential network changes that could mitigate disruptions. It will also enable assessment of impacts on dependent infrastructure and determine areas of common vulnerability. This report describes development of that capability.

Within the Communications Sector, the Internet provides vital services that enable many aspects of modern life—from driving the economy to providing life-saving emergency communications. When mature, this capability—coupled with the voice and emergency services capabilities—will allow integrated assessment of disruption impacts across the Communications Sector. Specifically, the Internet model will enable national and regional assessment of the impact of disruptions on Internet performance and connected infrastructure services such as Next Generation 9-1-1 communications. The model will also help to determine the viability of various disruption mitigation options. The capability can be used to identify common dependencies on communication paths and assets in infrastructure networks, such as banking and finance, which depend on the Internet for their operations. Scenario analyses will assist with preparation and planning and enable NISAC to study the effectiveness of mitigation options.

1.1. Overview

The Internet is a best-effort routed system and is self-correcting in the event of the loss of any particular route. Traffic may take alternate paths between the source and destination. There is no concept of an end-to-end dedicated connection for the duration of a “conversation.” Traffic is divided up into packets that are sent individually; each packet may take a different route based on the state of the network when that packet is sent.

The Internet relies on several key pieces of physical infrastructure to operate. These pieces of infrastructure are 1) the facilities that house networking and/or computing equipment; 2) long-haul and metro fiber cables; and 3) submarine cables.

The facilities that house equipment go by many different names, in part dependent on the services that are within them. Data centers are typically larger facilities that contain equipment for multiple entities. Services within them range from website hosting to large scale distribution and storage of data (e.g. cloud storage, streaming video, email serving) to interconnection of service providers. Many data centers also provide collocation services which allow third parties to rent space for computers and networking equipment.

Internet exchange points are interconnection facilities, typically within data centers, where Internet service providers can exchange traffic. These sites are owned and operated by various commercial entities, partnerships, and some universities. For example, an Internet exchange may rent a floor of a particular data center. Public exchanges are network or provider neutral meaning that any Internet service provider can establish a presence there. This does not guarantee that all providers at a given facility will exchange traffic with all other providers at that facility. Those agreements, called peering agreements, are separately brokered on a provider-by-provider basis. Private exchanges are sites where individual providers have made agreements to collocate equipment so they can exchange traffic. This may occur in data centers, telecommunications central offices, or at Internet points of presence.

Fiber optic cable is the physical connection that makes up what is thought of as the Internet backbone. Long-haul fiber is high capacity and runs between major interconnection points and major provider switching facilities to form the nationwide backbone. Metro area fiber is lower capacity and is used to connect facilities through a metropolitan area.

The final key piece of infrastructure is submarine cable. Transoceanic cables form the connections for the domestic Internet to the rest of the world. While there are a number of satellites that also perform this function, their capacity is significantly lower than that of the existing submarine cable infrastructure.

1.2. Organization of the Document

This document describes the construction and use of the nationwide Internet model, Athena. Section 3 discusses the model construction including data used and integrated to create the nodes in the network and the connections between them. Section 4 describes the use of the model to analyze network disruptions and mitigations. Section 5 discusses conclusions and potential next steps.

2. CONSTRUCTING THE MODEL

The model is designed to analyze the impacts of physical facility disruptions on the Internet. The base information used to construct the model is the Fiber Lit Buildings dataset from GeoTel.[1] The dataset contains records on Internet Service Providers (ISPs) which include street address, geospatial location, and facility category. Facility categories are broken down into central offices, data centers, and points of presence. An ISP at a particular location is a node in the network.

For the initial version of the model, ISPs were selected for inclusion based on the Center for Applied Internet Data Analysis (CAIDA) Autonomous System (AS) Rank information.[2] Autonomous Systems (ASes) can be roughly mapped to ISPs. For example, Level 3 Communications, maps to an AS named LEVEL3 and also an AS named LVLT-3549 among others. CAIDA ranks ASes by a variety of methods; we chose to use the ranking based on the providers with the largest number of direct and indirect customers, filtered to capture primarily US-based providers. Since provider names and relationships can change over time due to mergers and acquisitions or contractual changes, care was taken to use data of similar vintage from both CAIDA and GeoTel to determine rankings and relationships. The initial selection of providers includes the twelve largest US providers, plus two additional regional providers to better study the network interactions.

In some cases, ISP names did not directly map to ASes due to mergers and acquisitions, or an ISP operating as an older AS name. In order to translate company names from the Fiber Lit Buildings dataset to current ISP names that could be associated with an AS, NISAC used the American Registry for Internet Numbers (ARIN) coupled with news and financial reports.[3]

2.1. Assumptions

NISAC made several sets of simplifying assumptions to develop the model. The assumptions will be discussed in more detail in their respective sections. Some assumptions will be relaxed in future development efforts. For the Internet model, NISAC assumes:

- Relationships between ASes can be rolled up to the organization level and still apply.
- Observed carrier relationships do not vary across regions.
- AS relationships between large carriers change infrequently except in cases of mergers and acquisitions such that datasets from CAIDA and GeoTel from similar time periods can be paired to obtain the appropriate relationships for the providers.
- Fiber routes and routers and switches are not capacity constrained.

[1] http://www.geo-tel.com/, accessed July 2017.
[2] http://as-rank.caida.org/, accessed July 2017.
[3] https://www.arin.net/, accessed July 2017.

- Collocated providers can interconnect provided the appropriate agreements are in place, and these interconnections can be changed in a matter of hours, at most.
- ISPs connect their equipment in ring configurations for their internal networks.

2.2. Creating Model-Ready Data

Once the providers are selected, the model-ready datasets (links and nodes) are generated from the raw GeoTel data through a set of steps to consolidate like records and to filter out records which are not of interest to the model.

2.2.1. Creating the Nodes

As mentioned above, nodes in the network represent an ISP at a single physical location. The GeoTel Fiber Lit Buildings dataset is used to create those nodes by talking the following steps:

1. Consolidate all provider specific records at the same location into a single record.

   A specific facility (based on address) could have a facility category of one up to all three categories assigned to it (data center, central office, point of presence). This results in providers at that location also having records divided by facility category. Since the nodes in the model represent a provider at a specific physical location, these were consolidated.
2. Filter the consolidated provider set to capture the desired ISPs obtained from the association of ASes to providers based on AS rank.

2.2.2. Creating an Individual ISP's Network

Once the nodes are created, they must be joined together in some fashion to create a representative model of the Internet. We begin the process with an individual ISP's network or the internal links between nodes belonging to the same ISP.

We developed an algorithm to create the logical connections or links in an individual ISP's network. This represents how Internet traffic flows within a single ISP's network. We wanted to closely approximate the ring network topology that ISPs use to physically connect nodes in their network. To accomplish this we sort the nodes for a given ISP by their location, select the first node, connect that node to its two nearest nodes, and repeat this process until all the nodes in a given ISP are connected. We nominally compared several of our logical network representation of an ISPs internal network with physical network diagrams for that ISP and they were similar.

2.2.3. Constructing the Links Between ISPs

After creating the internal logical connections for each ISP in our network, we have to create the inter-ISP connections. ISPs will transit traffic off their own network onto another ISPs network at points where those providers connect, at Internet exchange points or other facilities where they have chosen to collocate for that purpose.

The Border Gateway Protocol (BGP) is a standardized routing protocol used to exchange traffic among ISP networks. BGP connections create path vectors for Internet traffic based on the business relationships between ISPs. These business relationships determine the type, direction, and amount of traffic that can be exchanged as shown in Figure 1. The three main business relationship types are peer, provider, and customer. A peer relationship is a bi-directional link, note that in Figure 1 there is no monetary exchange shown between providers with a peer relationship. These providers typically exchange traffic as equals. Provider and customer relationships are each directional links and labeled from the perspective of the individual provider.

For example, in Figure 1, from the perspective of ISP C, ISP B is a peer, so we label this as a bidirectional peer link between ISP C and ISP B. ISP F is a customer of ISP C, so we label the directional connection from ISP C to ISP F as a customer link. The link in the opposite direction from ISP F to ISP C is labeled as a provider link. Similarly ISP A is a provider to ISP C, so the directional link from ISP C to ISP A is labeled as a provider link and the directional link from ISP A to ISP C is a customer link.

Labeling the links and connecting the nodes in this way, allows us to enforce a set of rules that prevent an ISP from carrying traffic for which it has no monetary incentive.

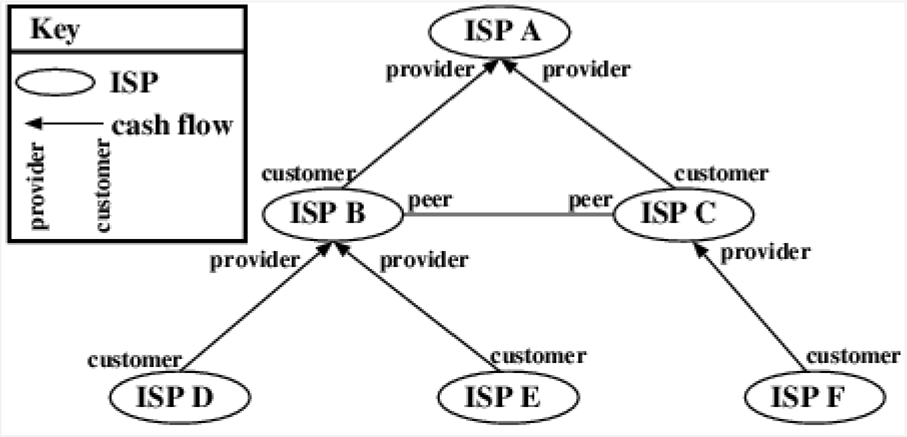

Figure 1. Types of Autonomous System Business Relationships[4]

The CAIDA AS relationship information also includes a relationship called siblings. This represents ASes that are administered by the same entity, usually as a result of mergers and acquisitions. We account for these relationships by rolling the relationships up to the largest administrative entity and applying those relationships across the associated ASes. A sample relationship table for the providers included in our model is shown in Table 1.

[4] http://as-rank.caida.org/, accessed July 2017.

To complete our network representation of the internet, we need to include the Internet exchange points and logically connect ISPs with one another based on their contractual relationships. We used the Center for Applied Internet Data Analysis (CAIDA) AS business relationships. To determine an Internet exchange point, we geo-located all the nodes to determine which ISPs were collocated. For each collocation point, we created a logical connection between the collocated ISPs based on CAIDA AS business relationships. ISPs which did not have a stated business relationship were not connected. A sample of the constructed relationship table is show in Table 1. The relationships are defined from the perspective of Provider 2. Thus row one of the table shows that Time Warner Cable is a customer of XO Communications. While row six shows that Cogent Communications is a provider to Lightower.

Table 1. Sample Relationship Table[5]

Provider 1	Provider 2	Relationship
XO COMMUNICATIONS	TIME WARNER CABLE	Customer
XO COMMUNICATIONS	VERIZON	Peer
XO COMMUNICATIONS	SPRINT	Peer
XO COMMUNICATIONS	AT&T	Peer
LIGHTOWER	LEVEL 3	Peer
LIGHTOWER	COGENT COMMUNICATIONS	Provider
LIGHTOWER	HURRICANE ELECTRIC	Provider

Our network now has four types of links: internal, peer, provider, and customer. To remain consistent with AS business relationship we have three exclusion rules for how traffic transits across the network. The rules are designed to keep an entity from carrying traffic when no one is paying it to carry that traffic.

1. A customer of an ISP will not transit traffic for that ISP across a peer connection. For example, in Figure 1, ISP B will not send traffic that it receives from ISP A to ISP C. B pays A to send traffic to A, A does not pay B, nor does C. Therefore B has no incentive to carry this traffic. Alternatively B would carry traffic from D to C, because it is paid by D to carry traffic for D.
2. A peer cannot transit to a provider. In Figure 1, ISP B will also not send traffic it receives from ISP C to ISP A. In this case, B pays A for traffic that it sends to A, and C does not pay B. Therefore B has no incentive to send traffic for C to A.
3. A customer cannot transit to a provider. In Figure 1, if we imagine ISP B with a second provider, ISP G, A and G cannot use B as their communication pathway if they have no other valid relationship.

[5] Constructed from AS Rank 2016-06-01 IPv4 dataset, http://as-rank.caida.org/, accessed July 2017.

2.3. Constructing a Solvable Network

Under disruption, the network as constructed required a custom path solver using a brute force node-node path search to determine if the network was fully functional. This approach was resource intensive. A simulation would take days to run.

To create a network that could be solved to determine what areas may be disconnected by a disruption, we modified our representation of the network by embedding the AS business relationships into the topology. For each node in our network we created a peer, provider, and customer transit node as shown in Figure 2. In the diagram, Comcast is a customer of Level 3 and each listing under Level 3 of peer, provider, and customer represent a transit node. We then connected each transit node to other providers transit nodes based on the business relationship rules described earlier. In the diagram the exclusion rules read as 1) a customer network cannot be used to transit to that customer's peer; 2) a peer's network cannot be used to transit to a peer's provider; and 3) a customer's network cannot be used to transit to that customer's provider (assuming it's a different provider otherwise we have created a routing loop). Internal networks are represented by simply connecting across the peer, provider, and customer transit nodes for links between the same provider. As shown in Figure 3, this allows the link to carry the information of what kind of network it came from across the network to connect appropriately to another provider on the other side.

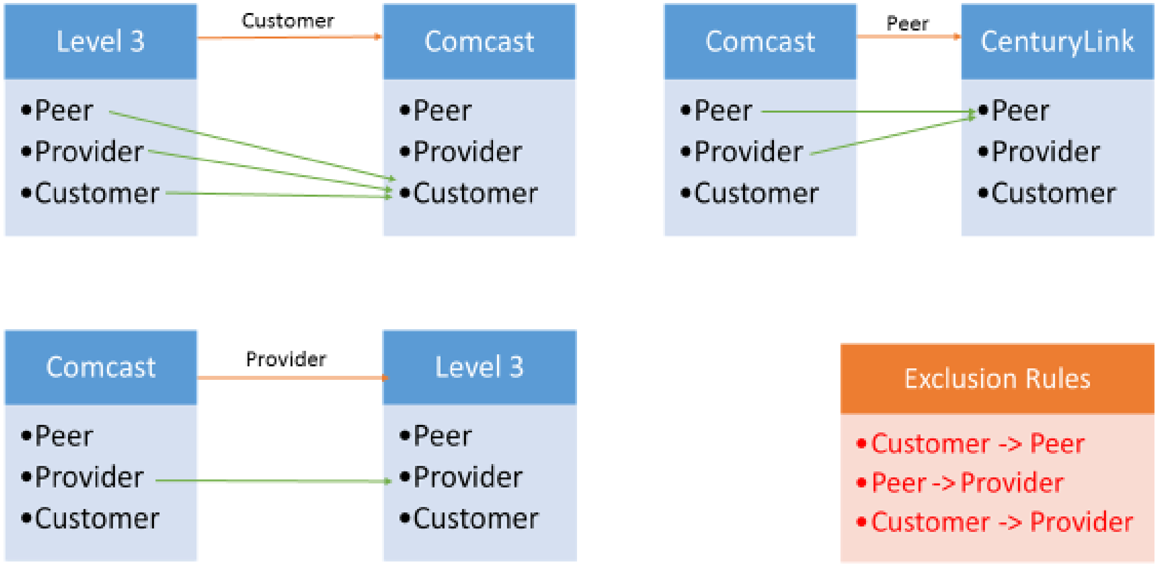

Figure 2. AS Relationships defined in the network topology

This approach eliminates the need for custom path algorithms that can accommodate path dependence (where you can go is governed by where you came from). Shortest path algorithms can determine the path between two nodes in the network since the topology carries the relationship constraints. We can also use algorithms that can determine whether or not the network has fragmented into more than one disjoint network, called clustering algorithms, to determine primary and secondary outages to the network.

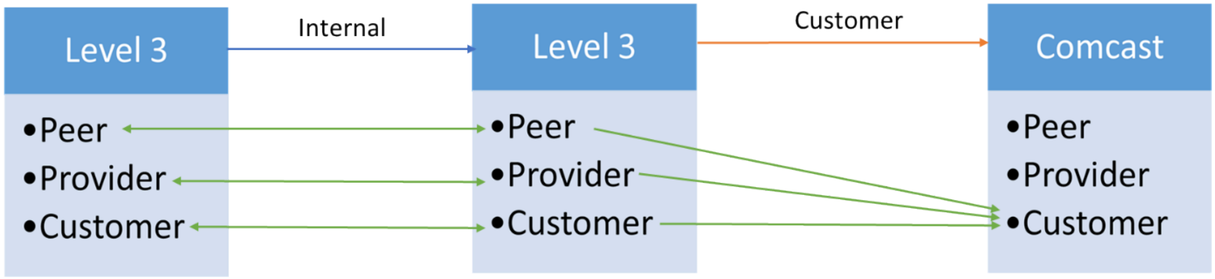

Figure 3. Internal Network connections

2.4. Service Territories

The model currently represents service territories by zip code. A zip code is considered out of service if every provider within that zip code is out of service. Otherwise, the model considers the zip code at risk for service outages.

3. MODELING DISRUPTIONS

The Athena disruption model uses graph theory to determine whether the current state of the model is a single connected network or a fractured network. Since we have represented the set of nodes and AS relationships topologically in our network we are able to use the connected component algorithm to determine the effect of a disruption to the network. Specifically we are using a breadth-first search connected component solver to determine how many sub-graphs our model contains. The algorithm traverses each node and link exactly once by walking through every node connected to the starting node, and then moving further away to explore the next level. A working network results in a single graph, or weakly connected component. A weakly connected component of a graph is a subgraph in which any two nodes are connected to each other by links. A disrupted network will result in multiple subgraphs that do not connect to one another.

To model disruptions, we first define what is out of service, or the primary disruption. Disruptions can be over a geographic extent, specific to a particular ISP to represent an ISP-wide outage, or simply to a set of nodes somewhere within the network. We remove the disrupted nodes and links from the model and apply a connected component algorithm to determine if there are any secondary network outages resulting from the primary disruption. If the connected component algorithm returns more than one weakly connected component, we know that there are nodes secondarily disrupted.

The following is an example of a disruption at a key node for traffic between a Zayo fiber-lit building and a Level 3 fiber-lit building. This example illustrates the importance of the AS relationships and their potential impact when nodes are disrupted. The base network with traffic flow between the Zayo node and the Level 3 node is shown in Figure 4 as the heavy gray dashed arrows. In Figure 4, Figure 5, and Figure 6 the red, blue and gray arrows represent peer, internal and customer relationships, respectively, and the directionality of those links represent normal traffic direction.

In the routing configuration shown, traffic flows from Zayo to its peer, Cogent Communications, traverses the internal network of Cogent and then to Cogent's peer, Level 3.

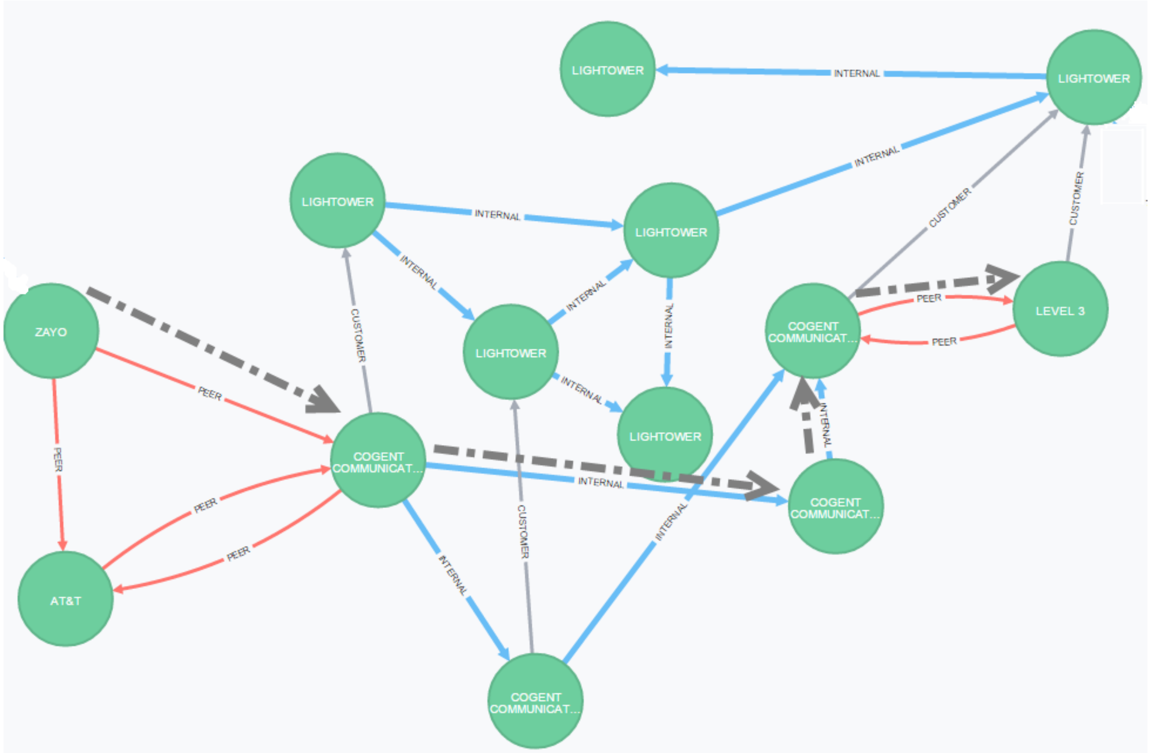

Figure 4. Typical BGP Routing and Traffic Flow

Figure 5 illustrates a disruption at a Cogent Communications fiber-lit building. With the connection between Cogent and Level 3 broken due to the outage, the only option would be for the traffic to transit the Lightower network to reach Level 3. Unfortunately since Lightower is a customer of both Level 3 and Cogent, that path is invalid since traffic would be using a customer as transit to that customer's provider. Traffic is not able to reroute to Level-3 due to AS relationships and thus Level-3 is considered a secondary disrupted node. The mitigation to this disruption is shown by use of the yellow highlighted links to route traffic in Figure 6. The heavy grey dashed arrows show the new traffic pathway, which traverses a red provider link from Lightower to Level 3 to allow traffic from Zayo to reach its destination. This pathway would normally not be available to traffic from Zayo to Level 3. By relaxing the AS relationship rules and allowing traffic from a Cogent peer link to transit Lightower to Lightower's provider, the Level 3 node is reconnected to the network.

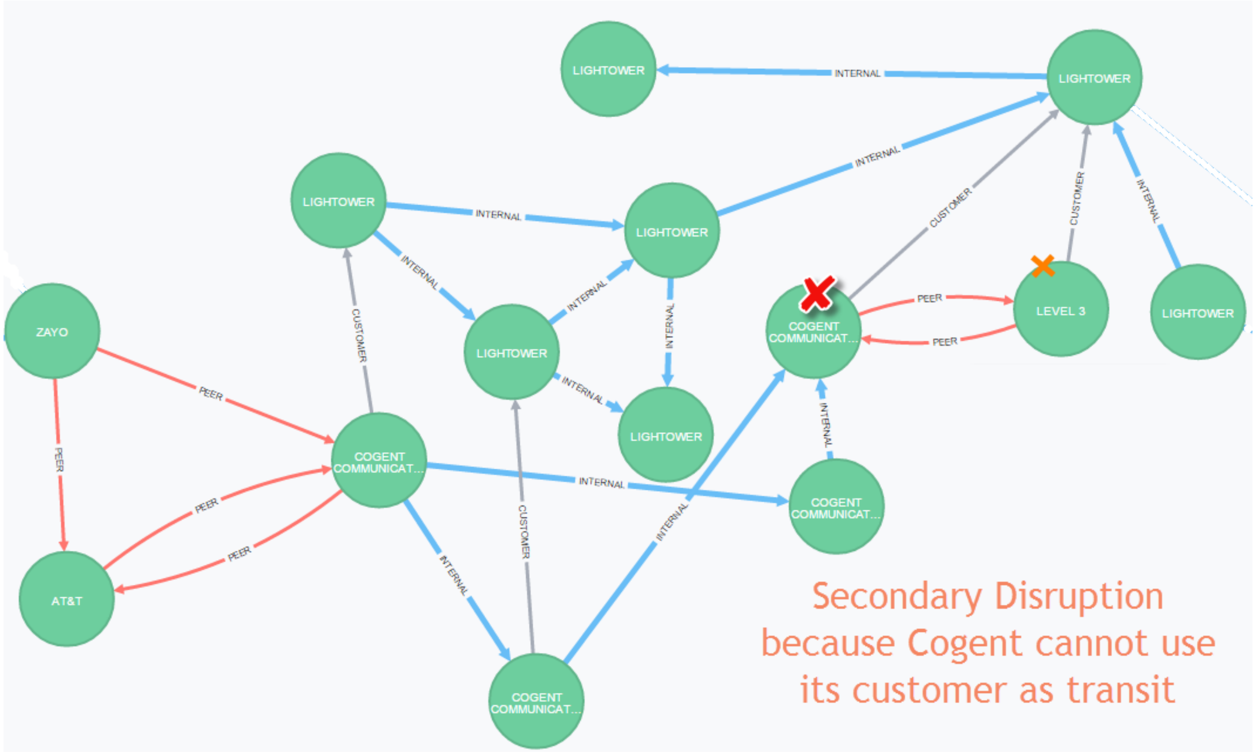

Figure 5. Secondary Disruption found with BGP Algorithm

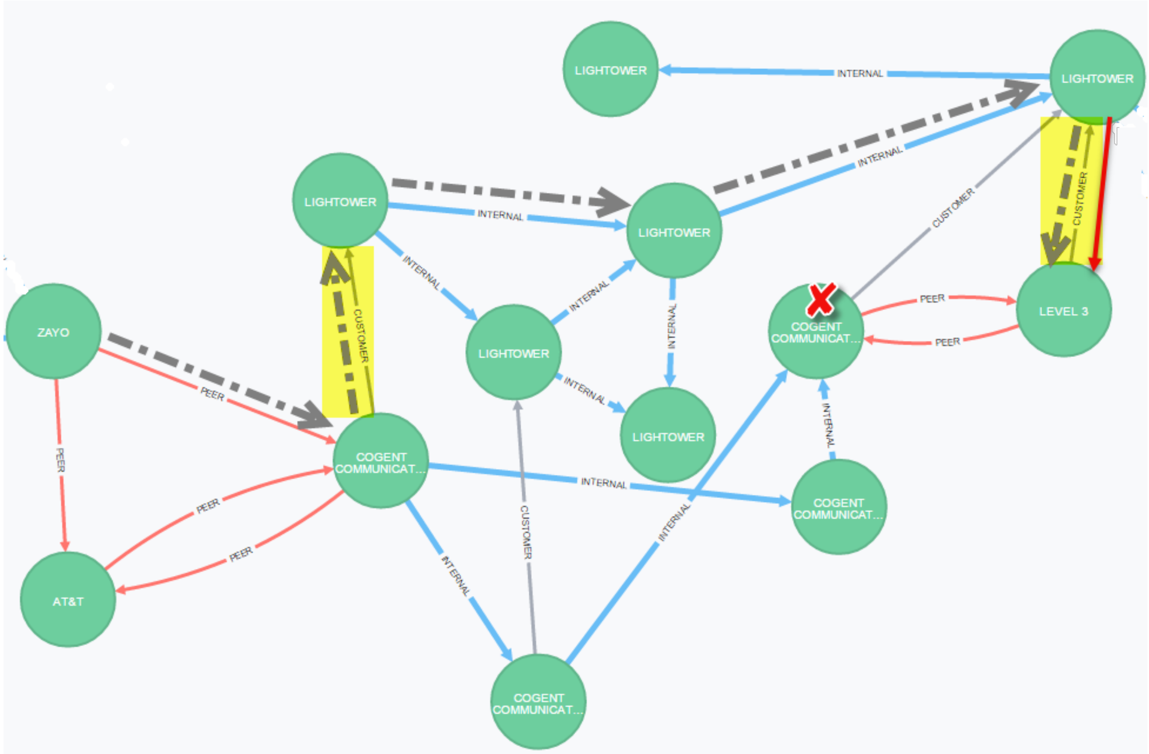

Figure 6. Mitigated Disruption by Modifying a Transit Rule

Additionally, we can also apply alternative network configurations to test network resiliency to changes in policy, such as the relaxing of IS relationships or additional relationships to the network.

4. SUMMARY AND NEXT STEPS

NISAC has developed a nationwide model of the Internet to study the potential impact of the loss of physical facilities on the network and on other infrastructures that depend on the Internet for services. The model looks at the Internet from the perspective of ISPs and their connectivity and can be used to determine how the network could be modified to assist in mitigating impacts of a disruptive event.

The network topology is based on the AS business relationships. This topology allows for the use of graph theory and associated algorithms to determine whether or not the network is fully functional or has been disconnected at some point(s). A working network results in a single graph, while a disrupted network results in multiple disconnected graphs. The analyst can then explore the graph to determine where the network could be reconnected by changing a business relationship rule if possible, or adding additional connections to the network.

Through the course of the work, we have identified several next steps to move the Internet modeling capability forward:

- Automating the process of network creation

  The current process to create networks could be automated to allow for easier changes to network configuration or addition of service providers.

- Creating a graphical user interface for analysis

  Currently the process for analyzing network disruptions requires an analyst familiar with network analysis tools based on graph theory. A graphical user interface could make the model more accessible to additional analysts.

- Expansion of the model to include additional regional providers

  Modifications to the network creation algorithms are needed to more accurately describe a regional providers network and constrain the network creation to specific regions. For example, several providers exist in very specific regions, so their networks would be ring networks in those regions; they then connect across either leased infrastructure or via agreements to use another provider's infrastructure, to link the disparate portions of their network together. The model's current internal network development algorithm does not account for disparate networks connected in this manner.

- Automation of the process to determine where to reconnect the network

  This is currently a manual operation and could be automated to allow for better analysis of options.

- Further refinement of service territories to beyond the zip code level if feasible

  Within the past year, the FCC has begun releasing residential fixed Internet connections by census tract rather than by zip code[6]. This information could allow

[6] https://www.fcc.gov/internet-access-services-reports, accessed July 2017.

for more fine-grained service territories and should be explored. In addition, since service territories for services provided by traditional wire line carriers are tied to wire center boundaries, there may be an opportunity to combine wire center service territories with the census tract information to better determine outage regions.

DISTRIBUTION

3 Department of Homeland Security Office of Cyber and Infrastructure Analysis
Attn: K. Kilby, R. Hanson, A. Blanning
NPPD/OCIA
245 Murray Lane, M/S 0390
Washington DC 20528-0390

1	MS1138	Nancy Brodsky	8873
1	MS1138	Rossitza Homan	8815
1	MS0899	Technical Library	9536 (electronic copy)

Sandia National Laboratories