Government Emergency Telecommunications Service (GETS) Survivability Analysis

presented by

Martin J. Fischer, Ph.D. —— Denise M. Bevilacqua Masi, Ph.D. —— David A. Garbin

Noblis

presented at the

13th International Command and Control Research and Technology Symposium (ICCRTS)
Bellevue, WA
June 17–19, 2008

13TH ICCRTS
Thirteenth International Command and Control Research and Technology Symposium

noblis
For the best of reasons

Acknowledgements

This work was partially funded by

the National Communications System Contract Number NBCH-D-02-0039 (Task Order Number 0200390095)

Research funding was also provided by the Center for Network-Based Systems, a partnership between Noblis and George Mason University

We would like to thank Dr. Pat McGregor of Nyquetek for reviewing our assumptions and giving us feedback

Outline

- Background and Purpose
- Survivability Background
- Approach and Scenarios
- Logical versus Physical Network Connectivity
- Shared Risk Groups
- Random Failures Analysis
- Directed Attack Analysis
- Natural Disaster Analysis
- Summary
- Next Steps

Background

Key question —
How to ensure effective emergency telecommunications during times of disaster when telecommunications resources are impaired and congested?

- The mission of the National Communications System (NCS) includes assisting the President, National Security Council, Office of Science and Technology Policy, and the Office of Management and Budget in —
  - The coordination of the planning for and provision of *national security and emergency preparedness communications* for the federal government under all circumstances, *including crisis or emergency, attack, recovery, and reconstitution*

Background (Concluded)

- In support of its mission, NCS runs several emergency telecommunications priority service programs for federal government users, such as the Government Emergency Telecommunications Service (GETS)

- However, industry is moving from circuit switched to Internet Protocol (IP) technology for all applications
  - NCS programs need to evolve toward IP capability to ensure continuity of priority traffic during emergencies—Next Generation Network (NGN) GETS

Purpose

The purpose of this presentation. . .

. . . is to provide analysis of survivability requirements for NGN GETS Voice given in the National Security/Emergency Preparedness (NS/EP) IMS Core Industry Requirements, Voice Service Document, to show that the requirements are defendable through modeling, or that the current requirements should be modified

Survivability Background

Survivability is defined as—

"A property of a system, subsystem, equipment, process, or procedure that provides a defined degree of assurance that the named entity will continue to function during and after a natural or man-made disturbance; *e.g.*, nuclear burst" *

—— also referred to as *Resilience* ——

*The current GETS survivability requirement is—***

NGN GETS Voice call/sessions should be able to be established and sustained between more than 65 percent of the pairs of surviving network entry points through which calls/sessions originate, and exit points through which calls/sessions terminate, when 10 percent of the nodes and links are randomly disabled.

* *Alliance for Telecommunications Industry Solutions (ATIS) Telecom Glossary 2000 (ATIS, 2000)*
** *NGN GETS Survivability briefing dated June 28, 2007(Survivability 20070618.PPT)*

Approach and Scenarios

- Use IP-SURVIV, a Noblis Survivability Analysis and Resilient Network Design Tool, to analyze throughput and connectivity of a network similar to a publicly available major ISP backbone as a function of outages in the network
- Failure scenarios for analysis
  - Random electronics failures
  - Directed/terrorist attack—affects major cities
  - Natural disaster—affects links and nodes in a region
- Metrics of interest
  - Percent of origin-to-destination (O-D) pairs communicating
  - Throughput—the amount of data per time unit that has a route available to its final destination*
  - Maximum round trip time

** Throughput and percent of O-D pairs communicating are the same if traffic is uniform between all O-D pairs*

Logical Versus Physical Network Connectivity

Backbone

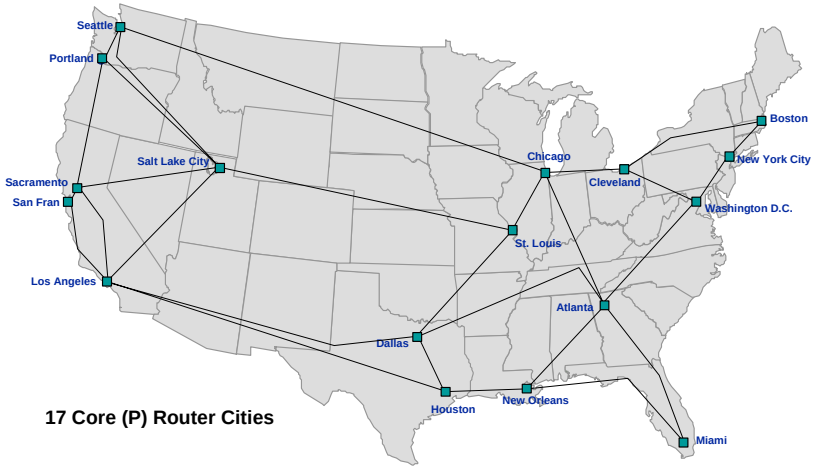

17 Core (P) Router Cities

Edge

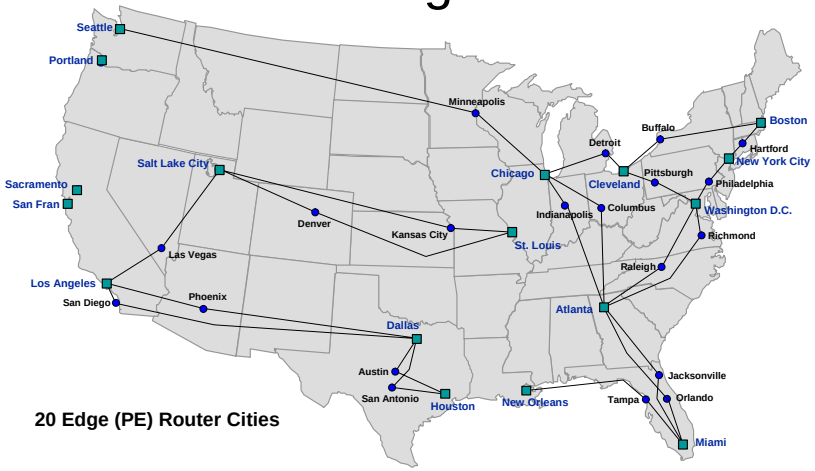

20 Edge (PE) Router Cities

Carrier Cities

37 Cities
Notional Fiber Connectivity

- Logical backbone and edge links in a network are physically routed over a limited number of optical fiber optic transmission paths
- The disruption of a single transmission segment can simultaneously fail several logical network links

Shared Risk Groups

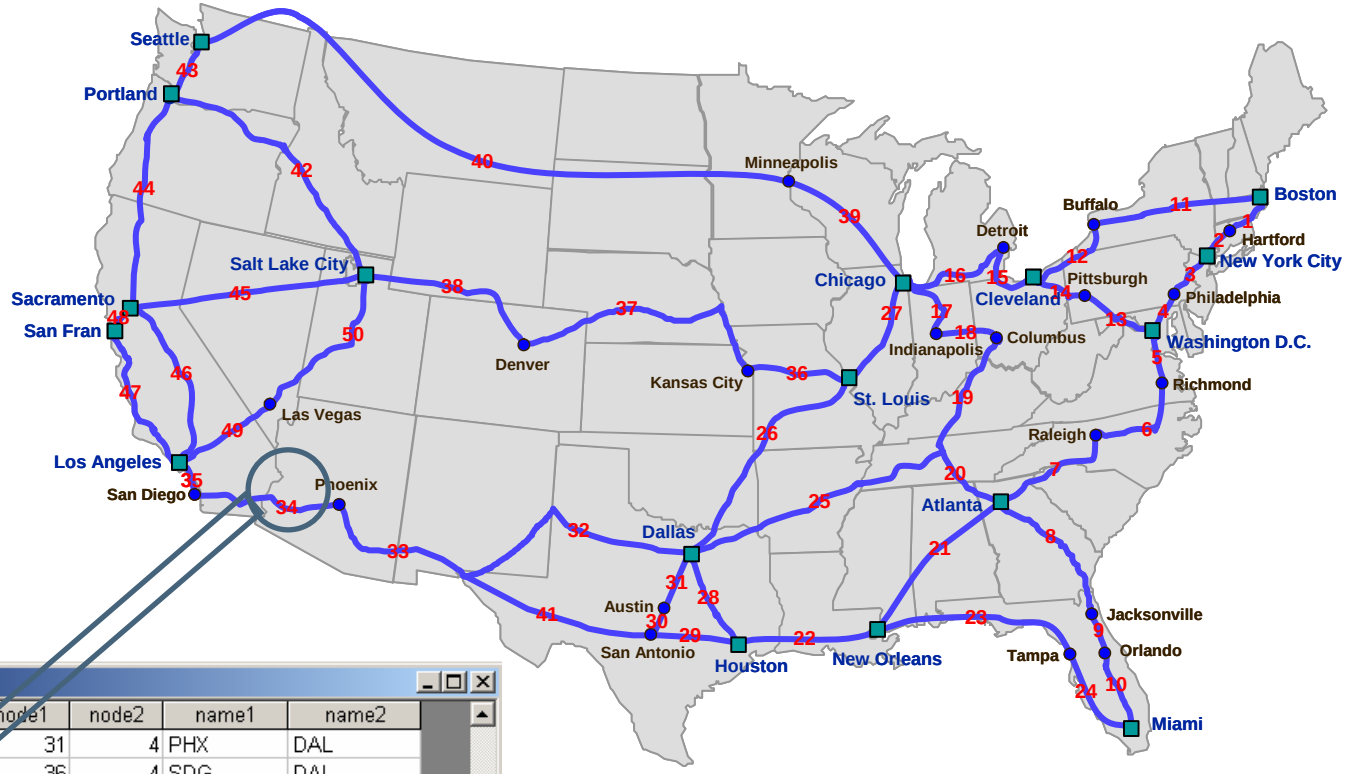

LINK GRP : Table

id	lnk_id	node1	node2	name1	name2
33	58	31	4	PHX	DAL
33	68	36	4	SDG	DAL
34	14	4	6	DAL	LAX
34	17	5	6	HOU	LAX
34	59	31	6	PHX	LAX
34	68	36	4	SDG	DAL
34	76	31	36	PHX	SDG
35	14	4	6	DAL	LAX
35	17	5	6	HOU	LAX

Record: 1 of 116

- Total number of fiber segments = 50
- Fiber segment 34 disruption causes outages on five separate network links

Random Failures Analysis

Random Failures Approach

- Throughput and connectivity of a network backbone representative of an major Internet Service Provider (ISP) backbone will be analyzed as increasing numbers of link groups in the network are randomly selected and disabled
  - Link groups for the network, which are sets of inter-switch connections lost from each segment outage, are first set up
  - The IP-SURVIV tool is configured to simulate disabling random selections of link groups

Random Failures Approach (Concluded)

- IP-SURVIV is run with the following assumptions
  - Open shortest path first (OSPF) routing
  - Population-based traffic (using 2005 Census data) between all node pairs, scaled so no packet loss in the baseline case
  - Segment outages are varied between 0 and 50 simultaneous segments disabled
  - Replications (n=10,000) are performed for each number of disabled segments investigated.
- Examine the throughput, connectivity, and Maximum Round Trip Time versus the percentage of transmission segments disabled

Throughput and Connectivity

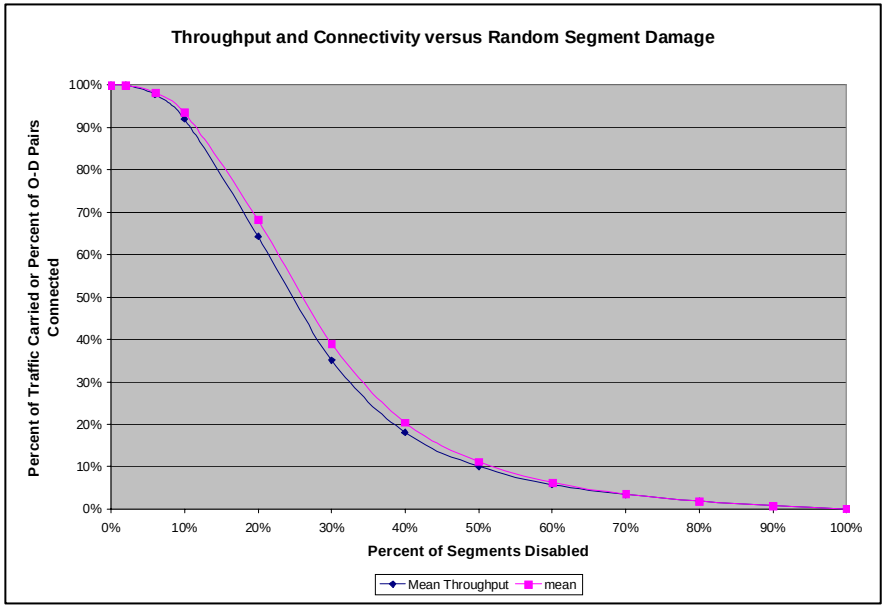

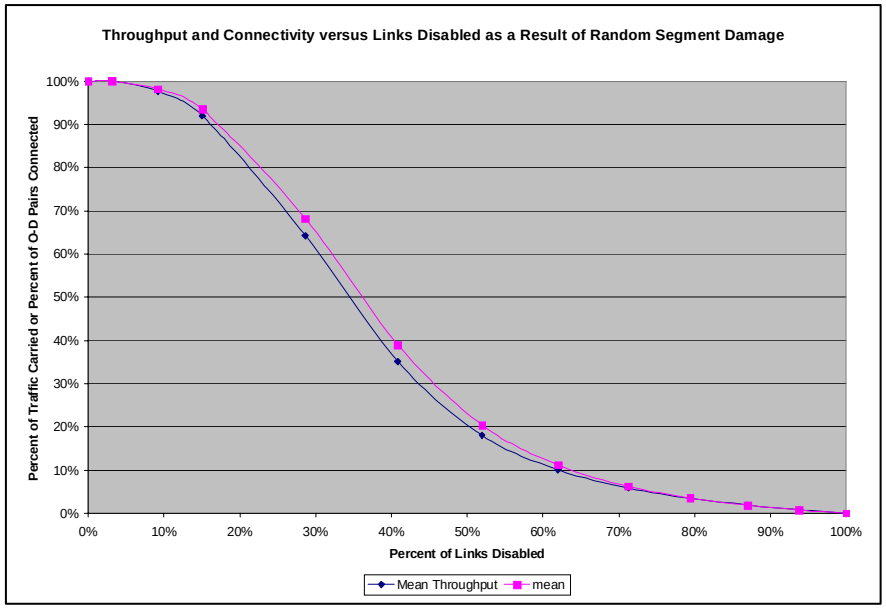

- Graphs show the throughput as a percent of traffic and percent of O-D pairs connected versus the percentage of segments and links disabled
- Throughput is slightly lower than connectivity due to traffic weighting
  - The impact of the traffic assumptions here are minimal; uniform versus non-uniform traffic results are close
- The percentage of links removed (right) are the links removed as a result of random segment outages (not random link outages)

Variability of Results for a Given Scenario

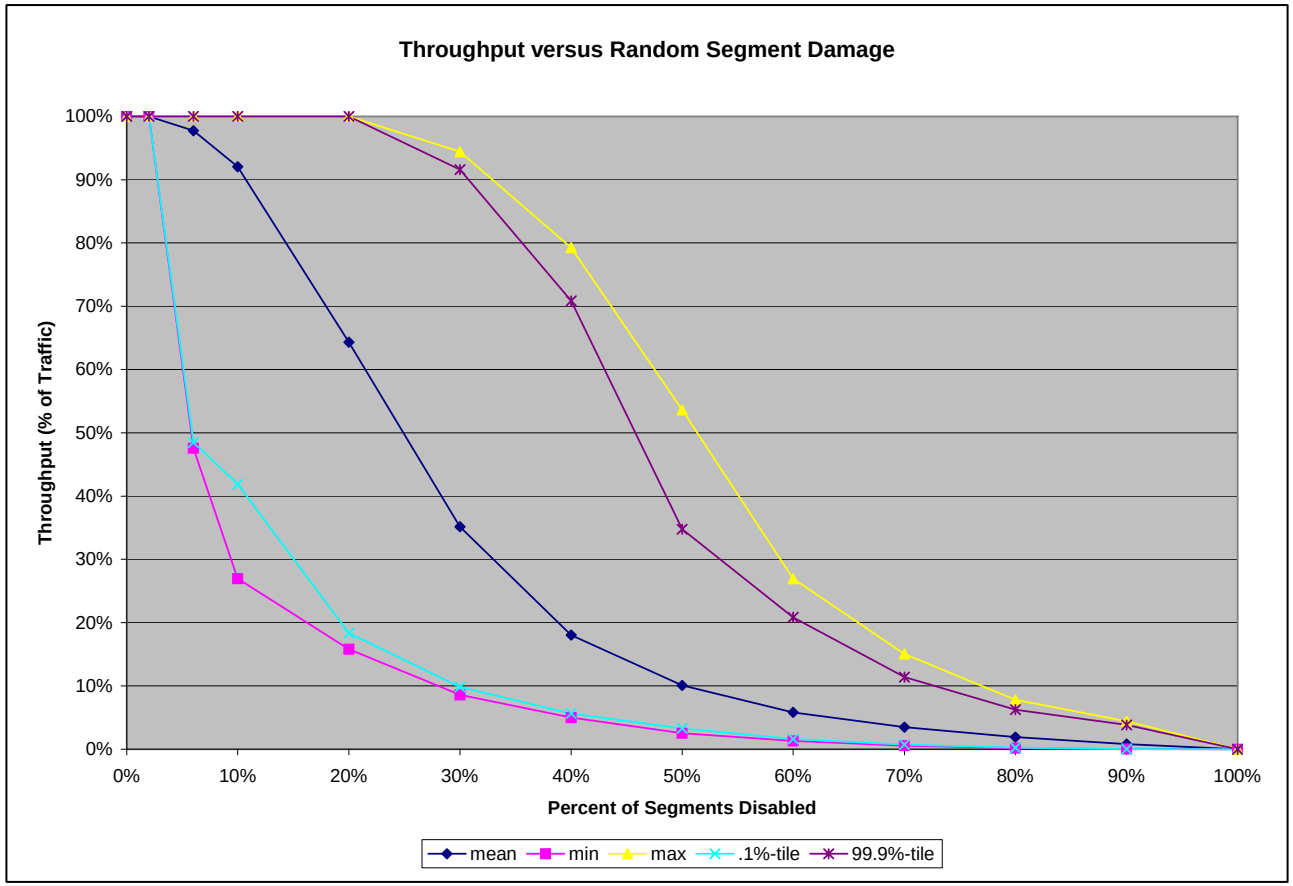

- There is much variability in the results for a given percent of segments disabled

Maximum Round Trip Time

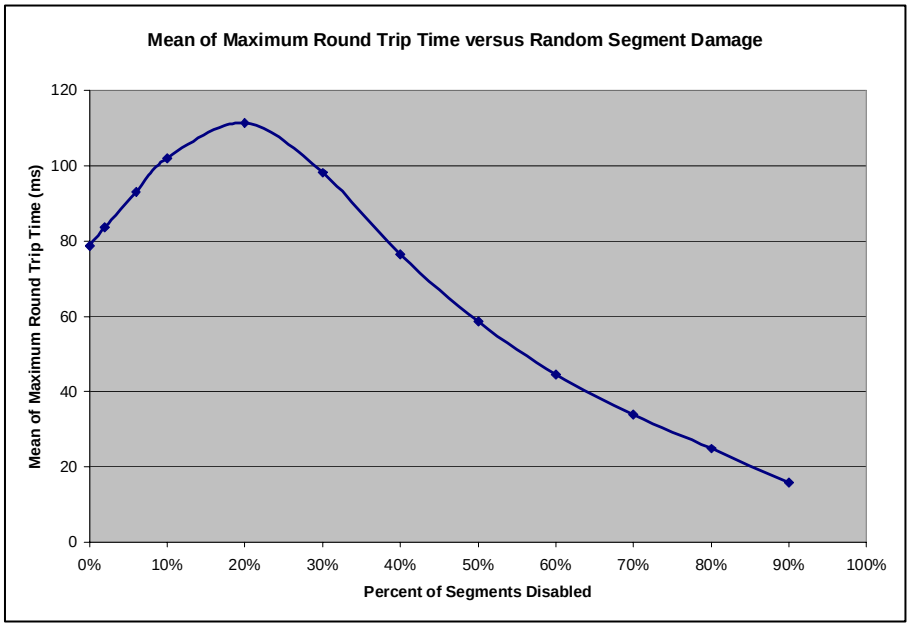

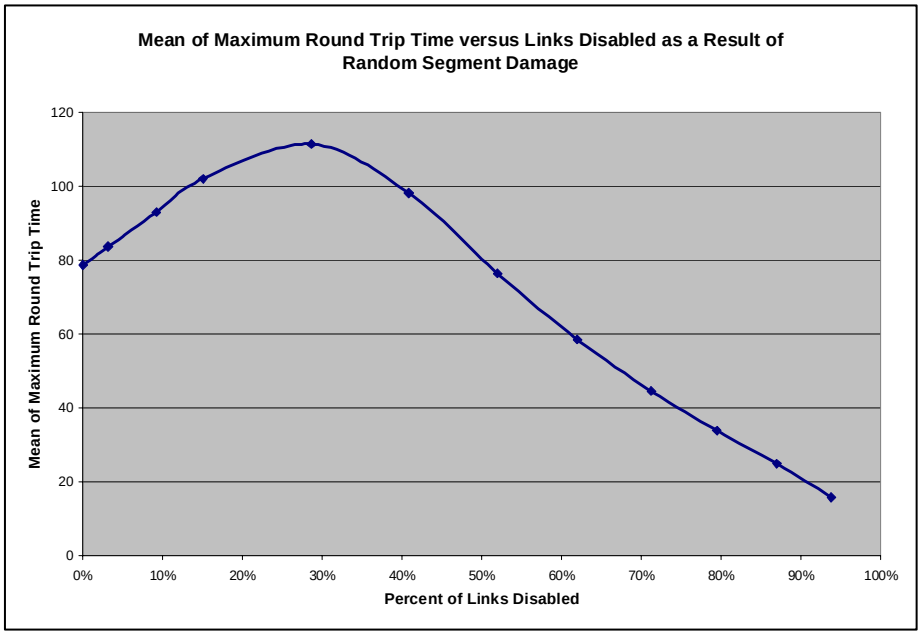

- Graphs show the mean of the Maximum Round Trip Time versus the percentage of segments and (non-random) links disabled
- Maximum Round Trip Time is greatest at 20 percent of segments disabled

Directed Attack Analysis

Directed Attack Approach

- Throughput and connectivity of a network backbone representative of an major ISP backbone is analyzed as network nodes are disabled, under the following scenarios
  - Single nodes outages
  - Multiple node outages, with several methods of selecting nodes for simultaneous node outage
    - Population based selection—top n population nodes disabled simultaneously (e.g., if two nodes are attacked, they will be the two highest population nodes)
    - Financial and government centers —New York, Los Angeles, Chicago, and Washington DC disabled simultaneously
    - Random selection—using n=10,000 replications for each number of disabled nodes investigated

Directed Attack Approach (Concluded)

- IP-SURVIV is run with the following additional assumptions
  - OSPF routing
  - Population-based traffic (from 2005 Census) between all node pairs

- Examine the throughput, connectivity, and Maximum Round Trip Time for different node outage scenarios

Single Node Outages

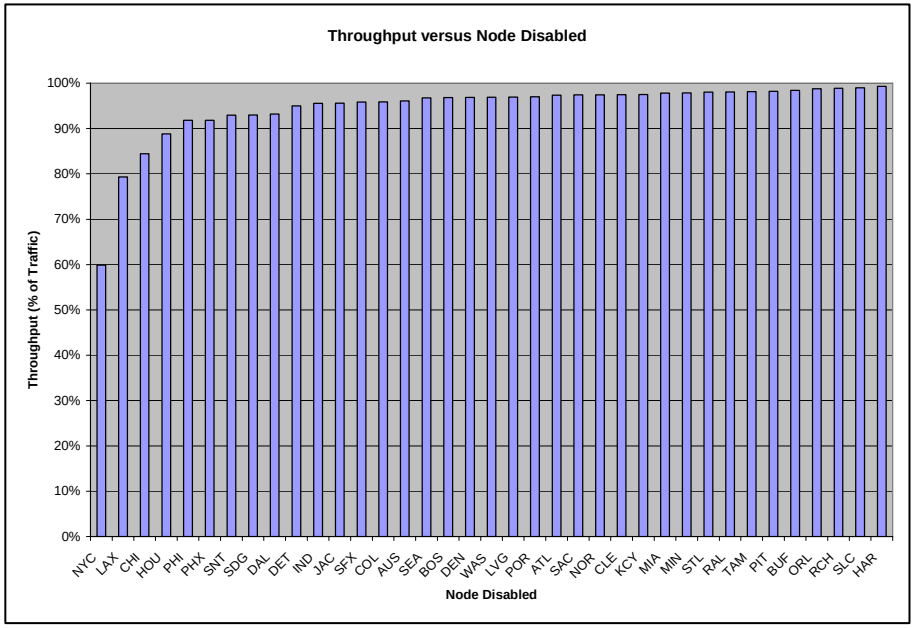

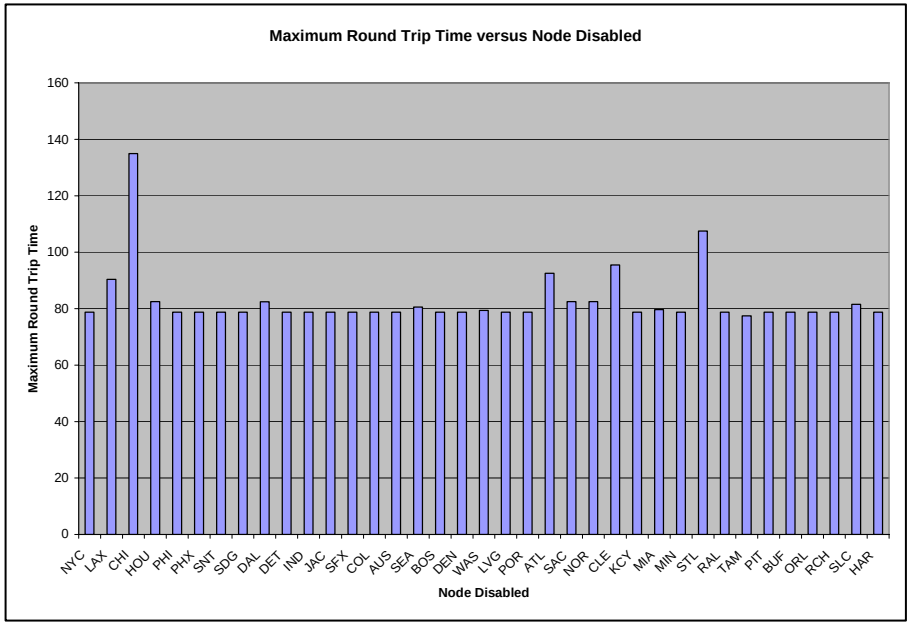

- X-axis in figures lists nodes in order of decreasing traffic level
- NYC, having the highest traffic levels, results in the greatest degradation in throughput if it is disabled
- Disabling centrally located nodes (CHI, STL) associated with links crossing the Mississippi River, has the greatest impact on Maximum Round Trip Time

Multiple Node Outages
Population Based

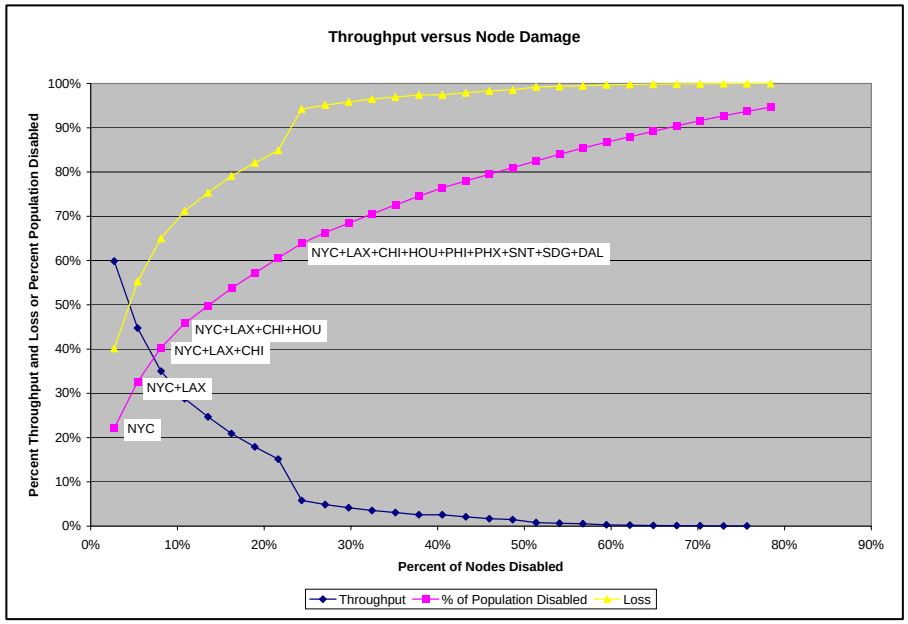

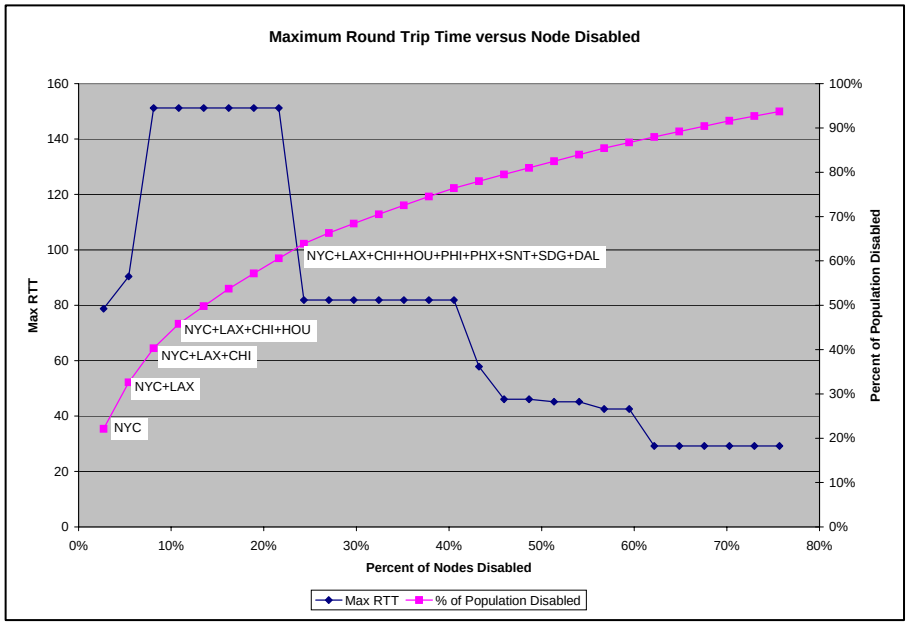

- Increasing numbers of nodes are disabled, in order of decreasing node population (and thus originating traffic)
  - For example, when three nodes are disabled, they are NYC, LAX, and CHI
- Traffic Loss increases (or Throughput decreases) at a higher rate than the amount of originating traffic or population that is disabled, due to other traffic that cannot reach its destination due to the outages
- Maximum Round Trip Time increases significantly when CHI is disabled in addition to other nodes
- Maximum Round Trip Time decreases for ≥ 24 percent of nodes disabled (the largest nine nodes), due to decreased traffic in the network

Multiple Node Outages
Financial and Government Centers

- Disabling New York, Los Angeles, Chicago, and Washington DC (financial and government centers) is compared to disabling New York, Los Angeles, Chicago, and Houston (highest traffic and population cities)
- Disabling New York, Los Angeles, Chicago, and Washington DC results in greater loss than if New York, Los Angeles, Chicago, and Houston were disabled, and a lower Maximum Round Trip Time

Nodes Disabled	Throughput (%)	Loss (%)	Maximum Round Trip Time	O-D Pairs Communicating (%)
NYC, LAX, CHI, WAS	21.9	78.1	126.2	51.1
NYC, LAX, CHI, HOU	28.8	71.2	151.2	79.3

Multiple Node Outages
Random

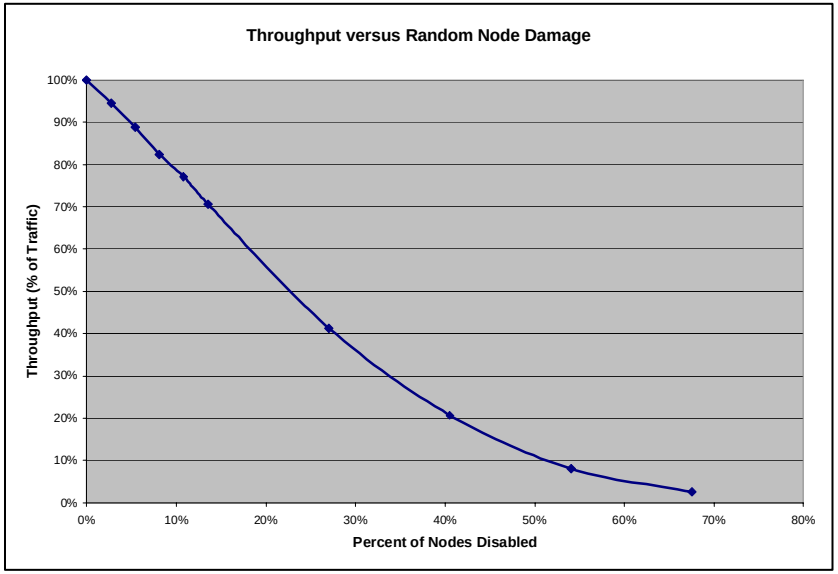

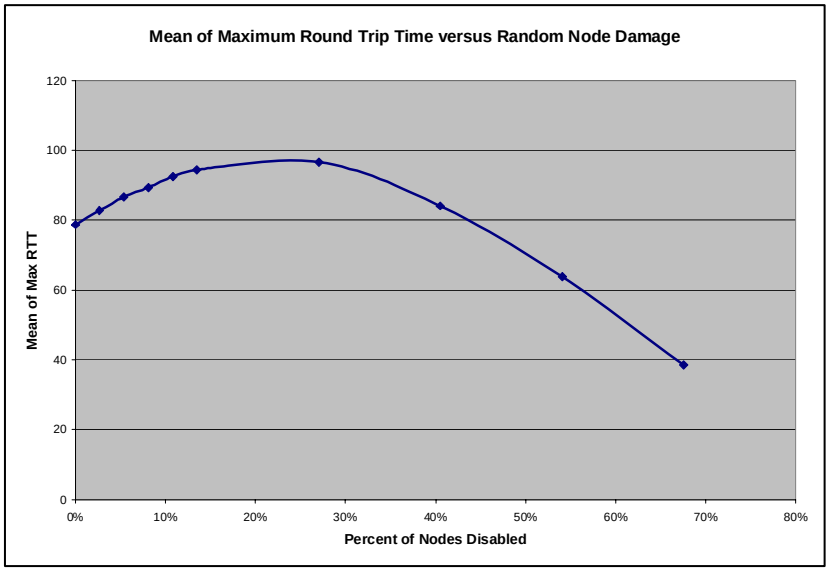

- If outages are selected in a random or less obvious manner, rather than based on financial, government, or population considerations, expected throughput could be much higher than in the previous slides
- For example, the average through-put when four nodes are randomly selected is much higher than for four large nodes (e.g., NYC, LAX, CHI, WAS)

Nodes Disabled	Throughput (%)	Loss (%)	Maximum Round Trip Time	O-D Pairs Communicating (%)
NYC, LAX, CHI, WAS	21.9	78.1	126.2	51.1
NYC, LAX, CHI, HOU	28.8	71.2	151.2	79.3
Mean Results for Four Randomly Selected Nodes	76.8	23.2	92.6	76.7

Multiple Node Outages
Random (Concluded)

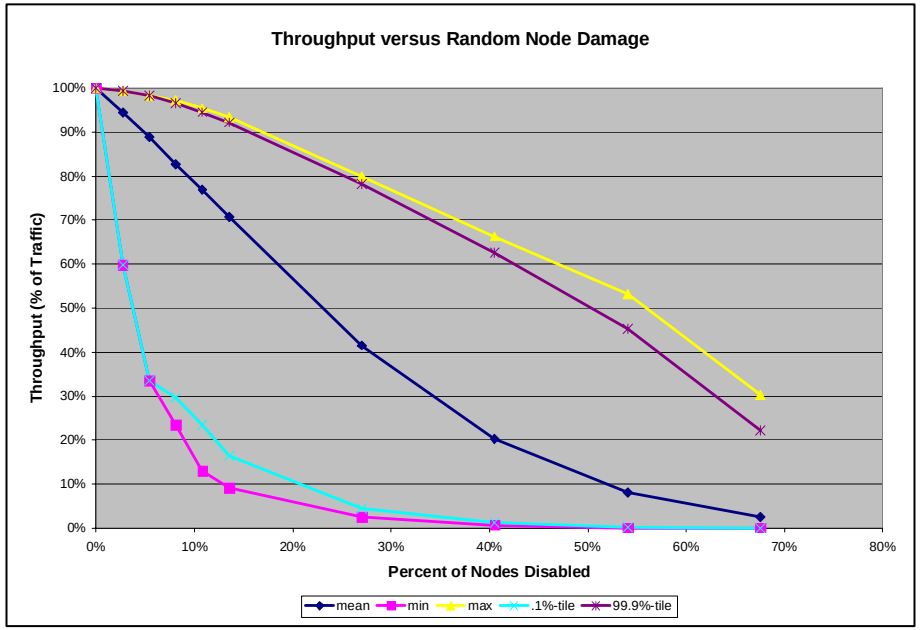

- As suggested by the previous slides, there is much variability in the results for a given number of disabled nodes
  - The combination of four nodes giving the lowest throughput (13 percent) in our simulation was Los Angeles, Chicago, Dallas, and Atlanta

Natural Disaster Analysis

Natural Disaster Approach

- Network survivability to natural disasters was investigated by completely disabling regions of the U.S., and running IP-SURVIV
- Same assumptions as other attacks (OSPF routing, population-based traffic)
- The U.S. Standard Federal Regions* were employed; outages were assumed to strike individual regions or neighboring groups of regions

** The ten standard Federal Regions were established by OMB Circular A-105, "Standard Federal Regions," in April, 1974. The regions are employed by a number of agencies including FEMA, HHS, GSA (GSA uses a slight deviation from the above). See http://www.fema.gov/about/contact/regions.shtm, http://en.wikipedia.org/wiki/Regions_of_the_United_States , or http://www.hhs.gov/about/regionmap.html*

Natural Disaster Impact

- The single regions which upon disabling most affect network throughput are
  - Region IX (6 nodes in CA, AZ, NV)
  - Region II (NYC node)
- When two regions are disabled, Regions V and VI and also Regions IV and V result in significantly reduced throughput (18 and 25 percent, respectively)
  - The disabling of either of these pairs of regions results in no routes for cross-Mississippi traffic

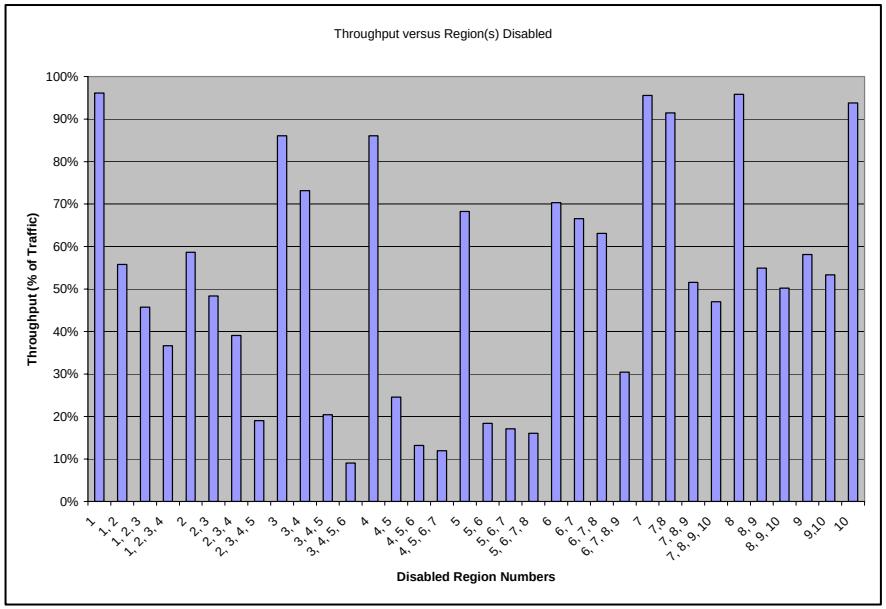

Natural Disaster Impact (Concluded)

- Disabling Region V has a greater impact on Maximum Round Trip Time than any other single region
  - Consistent with previous results of disabling the Chicago node alone
- When two regions are disabled, Maximum Round Trip Time is most affected for Regions VI and VII

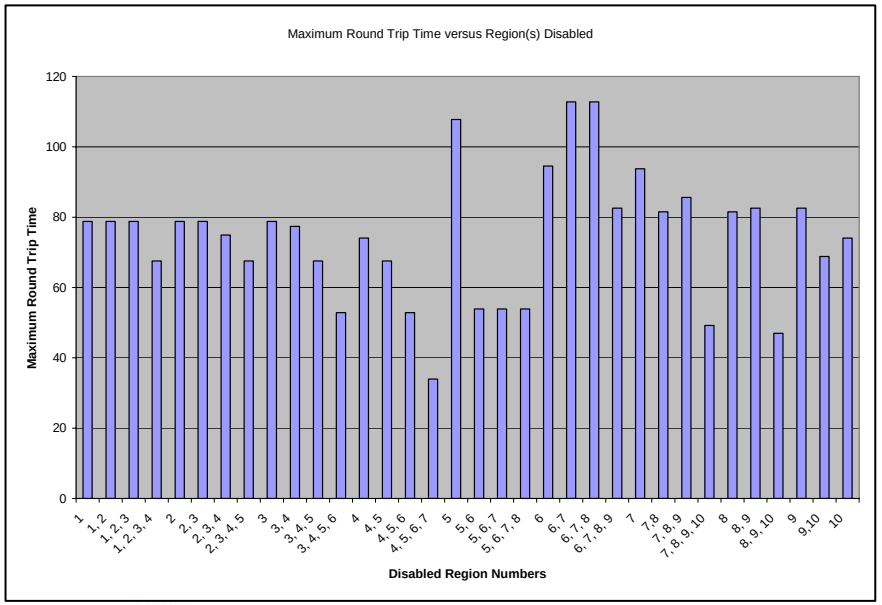

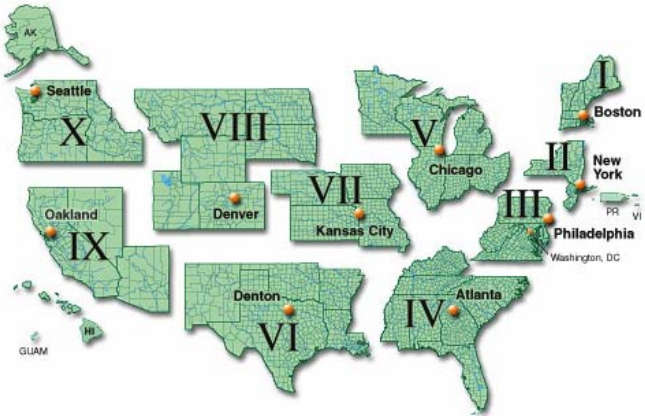

Summary

We have analyzed the impact of three types of failures on the survivability of an ISPs backbone network

- Random failures
- Directed attacks
- Natural disasters

Summary (Concluded)

- Analysis results will be used to modify survivability portion of the NGN GETS *Voice Industry Requirements* document
  - The current survivability requirement is not well defined, given the large amount of variability in the results
  - Due to high variability in the results, an objective, rather than a requirement, will be used during emergencies with high traffic overloads
- Throughput and percent of O-D pairs communicating are similar for average cases, indicating lack of sensitivity of results to the traffic matrix that was assumed
  - Throughput and percent of O-D pairs communicating are less similar for some low throughput cases

Next Steps

- Extend the current survivability analysis to multiple domains, and investigate whether there is a significant difference in the ability to communicate depending on whether or not the internetwork routing policies are modified to allow use of another carrier for transit
- Develop models to analyze GETS survivability with wireless access
- As both NGN and GETS survivability will depend on the ability of the call managers or processing elements to be reached by signaling, analyze the impact of the number, placement, and connectivity of call managers on the network survivability