SANDIA REPORT
SAND2014-18019
Unlimited Release
September 2014

Conceptual Framework for Developing Resilience Metrics for the Electricity, Oil, and Gas Sectors in the United States

Jean-Paul Watson, Ross Guttromson, Cesar Silva-Monroy, Robert Jeffers, Katherine Jones, James Ellison, Charles Rath, Jared Gearhart, Dean Jones, Tom Corbet, Charles Hanley, La Tonya Walker

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

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

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
US 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@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
US Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2014-18019
Unlimited Release
September 2014

Conceptual Framework for Developing Resilience Metrics for the Electricity, Oil, and Gas Sectors in the United States

Jean-Paul Watson, Sandia National Laboratories, Data Analysis & Informatics Dept.
Ross Guttromson, Sandia National Laboratories, Electric Power Systems Research Dept.
Cesar Silva-Monroy, Sandia National Laboratories, Electric Power Systems Research Dept.
Robert Jeffers, Sandia National Laboratories, Resilience & Regulatory Effects Dept.
Katherine Jones, Sandia National Laboratories, Operations Research & Knowledge Management Systems Dept.
James Ellison, Sandia National Laboratories, Electric Power Systems Research Dept.
Charles Rath, Sandia National Laboratories, Resilience & Regulatory Effects Dept.
Jared Gearhart, Sandia National Laboratories, Operations Research & Knowledge Management Systems Dept.
Dean Jones, Sandia National Laboratories, Operations Research & Knowledge Management Systems Dept.
Tom Corbett, Sandia National Laboratories, Resilience & Regulatory Effects Dept.
Charles Hanley, Sandia National Laboratories, Grid Modernization & Military Energy Systems Group
La Tonya Walker, Sandia National Laboratories, Earth Systems Analysis Dept.

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185

Abstract

This report has been written for the Department of Energy's Energy Policy and Systems Analysis Office to inform their writing of the Quadrennial Energy Review in the area of energy resilience. The topics of measuring and increasing energy resilience are addressed, including definitions, means of measuring, and analytic methodologies that can be used to make decisions for policy, infrastructure planning, and operations. A risk-based framework is presented which provides a standard definition of a resilience metric. Additionally, a process is identified which explains how the metrics can be applied. Research and development is articulated that will further accelerate the resilience of energy infrastructures.

ACKNOWLEDGMENTS

The authors of this paper wish to acknowledge the many experts that provided invaluable effort and information for this document. Dan Ton from the Department of Energy Office of Electricity Delivery and Energy Reliability and Rima Oueid and Kate Marks, from the Department of Energy Office of Energy Policy and Systems Analysis. Thanks is given for their tremendous leadership in formalizing the development of resilience concepts and bringing ideas to fruition. Acknowledgment is also given to Henry Willis from the Rand Corporation, who provided instrumental feedback and guidance. From the Sandia Corporation, thanks to Eric Vugrin for providing foundational help in establishing this document. Thanks to Lori Parrott and Steve Conrad for both having leveraged significant accomplishments via their long-standing contributions in NISAC. Further appreciation is given to Abe Ellis for his work in electric power resilience activities and Jason Stamp for providing leadership through the development of the Energy Surety Microgrid and design methodologies. And acknowledgement is also given to Bob Hwang for his continuous efforts and guidance.
We appreciate all of your efforts and look forward to collaborating in the future.

CONTENTS

Executive Summary 11
Conclusions and Recommendations 15

1. Introduction 17

2. What is Resilience? 19
2.1 Literature Review 19
2.2 Defining Resilience 19

3. Resilience Analysis Process 23
3.1 Define Resilience Goals 24
3.2 Define System and Resilience Metrics 24
3.3 Characterize Threats 24
3.4 Determine Level of Disruption 24
3.5 Define and Apply System Models 25
3.6 Calculate Consequence 25
3.7 Evaluate Resilience Improvements 25

4. Resilience Metrics 27
4.1 Metrics: Measures of Resilience 27
4.1.1 The Metric Is in Terms of Threat 27
4.1.2 The Metric Is Based on Performance 27
4.1.3 The Metric Measures Consequence 28
4.1.4 The Metric Accounts for Uncertainty 28
4.1.5 The Metric Effectively Captures Resilience 28
4.1.6 The Metric Is Not a Value Judgment 29
4.1.7 Multiple Metrics Are Often Necessary 29
4.1.8 System-Level Models in Resilience Metric Computation 29
4.2 Resilience Analysis Use Cases 30

5. A Framework for Developing Resilience Metrics 33

6. Populating Metrics 35

7. Simplified, Anecdotal Use Case 37
7.1 Define Resilience Goals 37
7.2 Define System and Resilience Metrics 37
7.3 Characterize Threats and Determine Level of Disruption 38
7.4 Define and Apply System Models 38
7.5 Calculate Consequence 39
7.6 Evaluate Resilience Improvements 39
7.7 Incorporating the Resilience Improvement Process 40

8. Advanced Use Case Summaries 41
8.1 Electricity 41
8.2 Petroleum 41
8.3 Natural Gas 42

9. What can be simplified today 43
9.1 How we can prepare for a more resilient framework 43

10. Research & Development 45
10.1 Furthering Resilience 45
10.1.1 Characterize Threats 45
10.1.2 Determine Level of Disruption 45
10.1.3 Define and Apply System Models 46
10.1.4 Calculate Consequence 46
10.1.5 Evaluate Resilience Improvements 46

11. Moving Forward 47
11.1 Identifying Specific Metrics that Could Be Generally Applied Across Many Different Systems 47
11.2 Stakeholder and Buy-In Process 47
11.3 Development of Sector-Specific System Models 49
11.4 Consequence Quantification 49
11.5 Conclusions and Recommendations 50

Appendix A: Reliability and Other Complementary Metrics 53
A.1 Resource Adequacy 54
A.1.1 Loss of Load Probability (LOLP) 54
A.1.2 Loss of Load Expectation (LOLE) 55
A.1.3 Expected Unserved Energy (EUE) or Loss of Energy Expectation (LOEE) 55
A.1.4 Effective Load Carrying Capability (ELCC) 55
A.2 Transmission Reliability Indices and Measures 57
A.2.1 Review of NERC Reliability Standards 57
A.3 Distribution Reliability Standards 62
A.3.1 Sustained Interruption Indices 62
A.4 Economic Perspectives on Electric Grid Reliability 63
A.5 Reliability Looking Forward 68
A.6 Integrated Reliability Index 68

Appendix B: Grid Resilience Metrics and Decision-Control Architectures 69

Appendix C: Electric Use Case 73
C.1 Illustrative Test System 73
C.2 Defining Threat Scenarios 74
C.3 Specifying Consequences of Loss of Delivery 75
C.4 Baseline Resilience Assessment 76
C.5 Restoration and Recovery Analysis 78
C.6 Investment Analysis 78
C.7 Advanced Planning 79
C.8 Summary 80

Appendix D: Petroleum Use Case 81
D.1 Scenario Description 81
D.2 Evaluating Resilience 82
D.2.1 Repair Model 82
D.2.2 Network Model 84
D.2.3 Consequence Model 86
D.3 Calculating the Consequence Metrics for the Scenario Earthquake 87

D.4 Evaluating Investment to Increase Resilience 90
D.5 Conclusions/Recommendations 91

Appendix E: Natural Gas Use Case 93

Appendix F: Methods to Normalize Resilience Metrics 99

Appendix G: Combining/Consolidating Resilience Metrics for a System 101

Distribution 103

FIGURES

Figure 1. A resilience metric framework is defined as "the probability of consequence X given threat Y". 12
Figure 2. Critical Infrastructure Planning and Operations Timelines: The top-most vector represents operator knowledge and actions. The bottom-most vector represents various aspects of system design and operation. 21
Figure 3. The Resilience Analysis Process: The steps of which progress counterclockwise f rom "Define Resilience Goals." 23
Figure 4. The resilience framework is a probability distribution of consequence to an extreme event(s). A resilience metric is a probability distribution of specific consequence to specific extreme event(s). 27
Figure 5. Infrastructure system attributes that affect the resilience metric. 29
Figure 6. Resilience Metric Framework 33
Figure 7. Example of a PDF describing resilience in terms of economic impact ($M USD). 35
Figure 8. Example of a PDF describing resilience in terms of lives at risk. 35
Figure 9. Performance indicators through time for one particular hurricane instance. 38
Figure 10. After transforming the performance indicators into consequence and accounting for uncertainty (left), the consequences for all realizations are plotted as a histogram of consequences (right). This histogram in converted into a probability density function, and is a resilience metric. 39
Figure 11. To analyze an investment alternative for resilience improvement, the resilience metric for the unimproved system is compared to that for the improvement alternative. 40
Figure 13. Incremental cost of reliability. 53
Figure 14. Hypothetical outage, system, and total cost as functions of reliability level. 64
Figure 15. Variation of reserve requirements with respect to customer outage costs. 67
Figure 15. Decision-control architecture for grid resilience. 70
Figure 16. IEEE 118 Bus Electrical Test System 73
Figure 17. Threat scenario tree representing three "children" nodes. 74
Figure 18. Detailed expansion of the hurricane threat. 75
Figure 19. Damage realizations for a Category 2 hurricane, landfall at high tide. 75
Figure 20. Example function relating MW not served at a bus to economic loss. 76
Figure 21. Resilience metric for the baseline system. 77
Figure 22. Resilience metrics for the same system under economic dispatch (left) and resilience dispatch (right) 77
Figure 23. Resilience metric focused on restoration and recovery. 78
Figure 24. Resilience metrics showing improvements over baseline by adding flood walls (left) and burying electrical cables (right). 79
Figure 25. Resilience metrics showing improvements over baseline by combining the addition of flood walls with buried electrical cables. The optimization algorithm selected the mix of each, constrained to $100M total. 80

Figure 26. The region impacted by an 1895 earthquake. The red region indicates the extent of severe damage to structures. The yellow region indicates the area over which shaking due this earthquake was felt by observers. 81
Figure 27. There are four major transmission pipelines that could be severely damaged by a New Madrid earthquake. Two of these, the Capline and the Midvalley carry crude oil. The TEPPCO and Centennial carry refined products. The shaded regions represent the shaking intensity of the scenario earthquake used for this analysis. Yellow and red indicate shaking intensities of VIII and IX, respectively, on the Modified Mercalli Intensity Scale. 82
Figure 28. Repair Cost trend calculated with the Repair Model. 83
Figure 29. Functional Infrastructure Capacity of damaged component calculated with the Repair Model. 84
Figure 30. The transportation fuel network model. 85
Figure 31. Simulated fuel consumption at terminals during 30-day disruption. 86
Figure 32. Assumed demand curve for fuel. 87
Figure 33. Probability distribution of repair times for each of for pipelines. 88
Figure 34. Histogram of 30 sampled repair times and resulting shortfall. 88
Figure 35. Histogram of 30 sampled repair times and resulting fuel cost increase. 89
Figure 36. Resilience metric for this analysis, shown as a histogram. 90
Figure 37. Improved resilience metric if two pipelines were relocated, shown as a histogram. 91
Figure 38. Map of the Southern California natural gas infrastructure in relationship to the San Andreas Fault. 93
Figure 39. Southern California network flow results summary for normal conditions (black), disrupted conditions with storage withdrawal restrictions (red), and disrupted conditions with no storage withdrawal restrictions (blue). 94
Figure 40. Histogram of economic impact shows baseline resilience of natural gas use case with a expected value of $325.1M. 96
Figure 41. Histogram of economic impact with unrestricted natural gas withdrawals shows natural gas use case modified resilience 96
Figure 42. Waiting on caption. 101

TABLES

Table 1. Mapping between goals, performance indicators, and their definitions 37
Table 2. A summary of the specific aspects of the resilience-enhancing process that each use case covers 41
Table 3: Example of capacity outage probability for a 6-generator system with FOR of 0.08 for each generator 54
Table 4: Direct Cost of a Summer Afternoon Outage with 1-Hour Notice 67
Table 5. Natural Gas Prices by Sector for California in December 2013 95

NOMENCLATURE

AHP	Analytic Hierarchy Process
DOE	US Department of Energy
DOE/EPSA	DOE Policy and Systems Analysis
DOE/OE	DOE Office of Electricity Delivery & Energy Reliability
EIA	US Energy Information Administration
FEMA	Federal Emergency Management Agency
FERC	Federal Energy Regulatory Commission
GPCM	Gas Pipeline Competition Model
HAZUS	Hazards US Multi-Hazard
IEEE	Institute of Electrical and Electronics Engineers
MMCF/D	million cubic feet per day
NERC	North American Electric Reliability Corporation
NGA	Natural Gas Association
NMSZ	New Madrid seismic zone
NTFM	National Transportation Fuels Model
PDF	probability density function
PPD	Presidential Policy Direction 21
PSLF	Positive Sequence Load Flow (GE simulation software)
PSS®E	Power System Simulator for Engineering (Siemens simulation software)
R&D	research and development
RAP	Resilience Analysis Process
SAIDI	System Average Interruption Duration Index
SAIFI	System Average Interruption Frequency Index
SCADA	supervisory control and data acquisition
SD	system dynamics
SMART	specific, measurable, attainable, relevant, and timely
SME	subject-matter expert
SNL	Sandia National Laboratories
VaR	value at risk

Page intentionally left blank.

EXECUTIVE SUMMARY

The President of the United States has called for a Quadrennial Energy Review (QER) and has appointed the Department of Energy's Energy Policy and Systems Analysis (DOE/EPSA) Office to lead an area on energy resilience. EPSA's role is to deliver unbiased energy analysis to the DOE's leadership on existing and prospective energy-related policies, focusing in part on integrative analysis of energy systems. This report is written to inform DOE/EPSA as they write the energy resilience section of the QER. It includes a general resilience metric framework and procedures for analyzing, quantifying, and planning for resilience of energy infrastructure systems. Additionally, the report provides use cases regarding electricity, petroleum, and natural gas to provide tangible examples of how these resilience metrics can be put into practical use.

Our nation's emergence as one of the world's most productive and innovative economies reflects broad access to abundant, reliable, and cheap energy. Unfortunately, the threats to our energy infrastructure, both natural and man-made, continue to grow. Recent events such as "Superstorm" Sandy and the 2011 Tohoku earthquake/tsunami in Japan caused major losses in human life and infrastructure, with prolonged recovery and financial impacts. Our nation faces significant risk from prolonged electrical outages. For example, problems with the power grid now cost the economy at least $150 billion per year[1] and have been trending upward.

In February 2013, the President strengthened and broadened our national position on critical infrastructure resilience by implementing Presidential Policy Directive 21 (PPD-21), *Critical Infrastructure Security and Resilience*. The directive applies to all critical infrastructures, but calls out energy infrastructures as being *uniquely* critical due to the enabling functions they provide across all other critical infrastructures. This document goes on to define resilience as "the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents."[2] <u>As an important step in avoiding conflict with well established reliability metrics and associated regulations, we have constrained our definition of resilience to high-consequence, low-probability events. Consequences, as stated in PPD-21, reflect social welfare. They go beyond the ability of a system to operate and address the vitality of our national safety, prosperity, and well-being.</u>

Measuring progress toward a more resilient energy infrastructure requires developing and deploying metrics that can be used to assess planning, operations, and policy changes for energy infrastructure. Metrics developed for this purpose are described in this report and are represented as probability density functions (PDFs) of consequences that may result from one or more threats to a system. This representation allows the analyst to understand the expected consequences

---

1 "The Smart Grid: An Introduction," Prepared for the US Dept. of Energy by Litos Strategic Communication, (http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/DOE_SG_Book_Single_Pages%281%29.pdf).

2 "Presidential Policy Directive—Critical Infrastructure Security and Resilience," White House press release, February 12, 2013, (http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil).

using its mean value, but also clearly identifies the range of possible consequences by viewing the shape of its distribution. Threats may include natural or man-made hazards such as hurricanes or physical threats. The red curve in the figure below is a notional representation of a resilience metric, herein denoted as a *resilience framework*. By specifying the applicable system, the threat, and consequences, the framework is transformed into a metric.

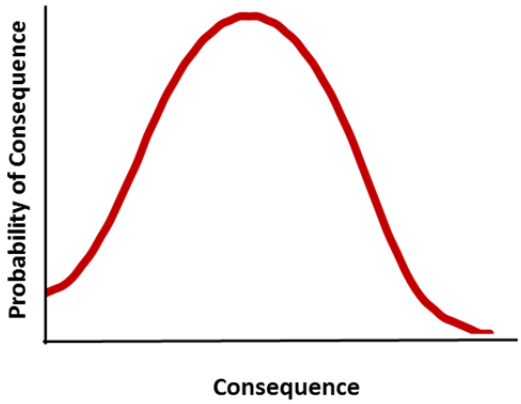

Figure 1. A resilience metric framework is defined as "the probability of consequence X given threat Y".

The framework allows for metrics that:

- **Are useful.** Metrics developed under this framework must be useful for decision making (by humans, computational analysis, or both). Decisions of interest include system planning decisions, real-time operations decisions, and policy decisions.
- **Provide a mechanism for comparison.** Applying the same metric to different systems should result in valuable information. Furthermore, the same metric must be able to differentiate between the resilience of a system that has not been enhanced and one that has (either through infrastructure or operations enhancements).
- **Are useable in operations and planning contexts.** The same metric should be able assist decisions for both planning and operations.
- **Exhibit extensibility.** The metrics selected must be scalable in time and geography. The metrics should remain valid as technology progresses and more complex analytic methods become feasible.
- **Are quantitative.** The framework must allow the development of metrics that can be used both qualitatively and quantitatively.
- **Reflect uncertainty.** It's critical that metrics are populated using methods that will quantify the uncertainty of the result. Specifically, decisions being made based on a resilience metric value must be well informed by the certainty of that value.
- **Support a risk-based approach.** The metrics should reflect a specific threat or set of threats, the system vulnerability, and potential consequences to people (beyond the immediate system effects).
- **Consider recovery time.** Resilience metrics should reflect the consequences over time, and therefore must consider the recover period either directly or indirectly.

Two fundamental concepts of this risk-based framework are

- resilience is defined with respect to disturbance(s) or threat(s) and
- consequences relate to the social effects of system performance in addition to system performance itself.

Although ongoing research and development (R&D) is needed to improve the ability to link energy system performance to social consequences, is is useful understanding what can be harnessed now.

In order to populate (or enumerate) the metric, analysis must be conducted. This analysis can be complex, and although this report provides examples, this is also an area for further research. Because different regions of the country have different potential hazards, stakeholders must select metrics that will help them measure progress toward their specific resilience goals. For example, the East Coast is at risk for hurricanes, although the desert Southwest is not. Similarly, the most important social consequences for different regions within the US would likely be dissimilar. Although the framework applies broadly, individual metrics should be developed with stakeholder feedback.

Deploying metrics in the form discussed represents a fundamental change in approach for defining energy system resilience within the energy industry and state and local governments. There has been little work that quantitatively expresses values of resilience. Much of the previous work has focused on defining system attributes that result in increased resilience, such as the number of critical spare parts in inventory, but has been unable to quantify the resilience benefit. Although attribute-based resilience metrics will continue to be valuable, this shift moves toward a quantitative, risk-based assessment useful for complex decisions. Such a shift requires stakeholder education, especially in metric application (see Appendices B, C, and D for example case studies). Further, it relies upon new data sets, the use of system models, and rigorous analysis. A comprehensive survey of resilience metrics[3] can be found in a report conducted by the RAND Corporation for the DOE. Other significant reports on resilience are also cited.[4,5,6]

A Resilience Analysis Process (RAP) has been developed herein as method for the assessment of baseline resilience and evaluation of resilience improvements. The RAP is designed to support decision makers high-level goals with a defensible, risk-based decision. The first six steps of the RAP give decision makers and stakeholders a method for assessing a system's baseline performance. When all seven steps are followed, the focus of the RAP expands identifying

---

[3] H. H. Willis and K. Loa, "Measuring the Resilience of Energy Distribution Systems," RAND Justice, Infrastructure, and Environment, PR-1293-DOE, July, 2014.

[4] "Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation (SREX): Special Report of the Intergovernmental Panel on Climate Change," C. Field, V. Barros, T. Stocker, et al. Eds., (Cambridge University Press, New York, N.Y., 2011), 594 pp.

[5] North American Reliability Corporation, "North American Reliability Corporation (NERC) Cyber Attack Task Force Final Report," May, 2012.

[6] B. Plumer, "Bad news: The U.S. power grid is getting pricier, less reliable," *The Washington Post* (Blogs), March 8, 2013, (http://www.washingtonpost.com/blogs/wonkblog/wp/2013/03/08/surprise-the-u-s-power-grid-is-getting-pricier-less-reliable/).

improvements that increase resilience. These improvements could be identified by analyzing or by optimizing the characteristics of these proposals to identify the best improvement strategies. A summary of the seven-step RAP is as follows:

1. **Define Resilience Goals**. Before determining the scope of the system relevant for analyzing and selecting appropriate metrics, it is essential to define high-level resilience goals. The goal set during this first RAP step lays the foundation for all following steps.
2. **Define System and Resilience Metrics.** The system under consideration and the resilience metric definitions determine the analysis' scope. This could include identifying a larger system's geographic boundaries, relevant time periods, and/or relevant components.
3. **Characterize Threats.** Threat characterization is critical to understanding how capable the system must be to absorb and adapt to different types of attacks or natural events. When evaluating resilience against multiple hazards, information about (1) the likelihood of each possible threat scenario and (2) the capabilities or strength of the threat are extremely important. In risk analysis, threat and consequence are used to understand which vulnerabilities are most important to address to reduce the consequences associated with the threat.
4. **Determine Level of Disruption.** Once an understanding of the relevant threats has been solidified, the attributes of each threat are used to determine the amount of damage to the system (infrastructure, equipment, etc.) that is likely to result from that set of threats. This is the RAP step where expectations about structural damage or other system impacts that could affect performance are defined.
5. **Define and Apply System Models**. The damage states outlined in Step 4 can then be used as input to system models—tying damage to system output levels. For example, anticipated physical damage (or a range of damage outcomes incorporating uncertainty) to an electric grid from an earthquake can be used as input to a system model that ties those outages due to damage to load not served within the system over time. Multiple system models may be required to capture all of the relevant aspects of the complete system. Furthermore, dependencies may exist between models.
6. **Calculate Consequence.** When evaluating resilience, direct impacts to system output as a result of damage are only part of the story. Most energy systems provide energy some larger social purpose (e.g., transportation, health care, manufacturing, economic gain). During this step, outputs from system models are converted to the resilience metrics that were defined during Step 2. When uncertainty is included in the RAP, probability distributions will characterize the resilience-metric values.
7. **Evaluate Resilience Improvements.** Unless the RAP is being undertaken purely for assessment purposes, it is likely that some decision or decisions must be made about how to modify operational decisions or plan investments to improve resilience. After completing a baseline RAP through the preceding steps, it is possible and desirable to populate the metrics for a system configuration that is in some way different from the baseline in order to compare which configuration would provide better resilience. This could be
   - a physical change (e.g., adding a redundant power line);
   - a policy change (e.g., allowing the use of stored gas reserves during a disruption); or

- a procedural change (e.g., turning on or off equipment in advance of a storm).

Conclusions and Recommendations

- A framework for energy resilience metrics has been created such that:
  - Energy resilience metrics quantify the expected consequence due to events that have low probability but potentially high consequence. Consequences focus on social welfare, extending beyond system impacts.
  - The resilience metrics rely on the performance of the system, as opposed to attributes of that system.
  - The resilience metrics incorporate the uncertainty associated with limited information about the system and the threat.
  - Resilience metrics quantify performance given uncertainty, providing insights into risk management and cost/benefit processes for planning, operations, and policy building.
- A resilience analysis process has been created that explains how to use resilience metrics. The process is flexible enough for use by different stakeholders and infrastructures. Stakeholder goals should drive the selection of metrics used for an analysis within the framework provided.
- Continued research is essential:
  - More research is needed to improve quantification of human/societal consequences based on reduced system performance in a disruption. Key areas for R&D investment include multi-category uncertainty quantification, modeling and simulation of disruption, recovery and repair, and adaptive system operation algorithms.
  - Developing a library of suggested performance indicators and recommended methods for translating those system outputs to common consequence measures is a necessary national research and development pursuit.
  - Data availability will be a challenge in the early stages of adopting these methods, so some effort is likely to be needed with respect to data collection and establishing associated best practices.
- Outreach and collaboration is necessary to define the types of decisions that will use resilience metrics, as well as the metrics' units of consequence, selection of threats, and quantification of uncertainty.
- A stakeholder group should be created for the refinement and standardization of metrics for electricity, petroleum, and natural gas sectors for the validation of this resilience metric framework. Specific areas that should be addressed include:
  - Differentiate reliability metrics from resilience metrics with input from state, federal and regional regulatory authorities and other stakeholders
  - Determine federal, state, and local government roles
  - Work toward stakeholder buy-in and coordination: federal and state regulators, utilities, asset owners, and other key stakeholders

- Conduct an expanded case study using data from one or more major utilities (in coordination with that utility)

1. INTRODUCTION

Over the past decade, resilience has emerged as a new design and operations goal for the critical infrastructure protection community. Historically, infrastructure security activities have primarily focused on preventing disruptive events. However, the critical-infrastructure community now recognizes that it is simply not possible to prevent all threats, at all times, for all infrastructure assets. Hence, critical infrastructure resilience—an infrastructure's "ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions,"[7]—is now recognized as a complementary criterion to security activities. Federal policy toward infrastructure resilience has been formalized in Presidential Policy Direction 21 (PPD-21), *Critical Infrastructure Security and Resilience*, which has designated the energy sector as "uniquely critical."

The Department of Energy (DOE) has further recognized the need for increased resilience in the energy sector. *Economic Benefits of Increasing Electric Grid Resilience to Weather Outages*[8] estimates the average annual cost of weather-related power outages to be between $18 and $33 billion over the past decade. *U.S. Energy Sector Vulnerabilities to Climate Change and Extreme Weather*[9] notes that the frequency of extreme weather events is expected to increase as a result of climate change, increasing risk not only to the electric power sector, but also to petroleum and natural-gas infrastructure systems.

This report recommends a combination of technologies, policies, information, and stakeholder engagement to strengthen the energy sector's resilience to climate change, and the *DOE's Climate Change Adaptation Plan*[10] affirms this position, making "improve the climate resiliency of all DOE sites" one of the plan's top four goals. In addition, as part of a White House initiative under EO 13653 *"Preparing the United States for the Impacts of Climate Change,"* DOE and DHS co-chair an interagency infrastructure resilience working group. DOE's commitment to resilience extends beyond extreme weather and climate change disruptions to include intentional acts such as cyber-attacks. *The Roadmap to Achieve Energy Delivery Systems Cybersecurity*[11] designates "resilient energy delivery systems are designed, installed, operated, and maintained to survive a cyber incident while sustaining critical functions" as the vision for the energy sector. Furthermore, DOE's new Operational Energy and Resilience Program is expected to contribute to both the development of new resilience technologies and to coordination with the Federal

---

7 *Presidential Policy Directive 21, Critical Infrastructure Security and Resilience*, February 2013, (http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil).

8 *Economic Benefits of Increasing Electric Grid Resilience to Weather Outages,* August 2013, (http://energy.gov/downloads/economic-benefits-increasing-electric-grid-resilience-weather-outages).

9 *U.S. Energy Sector Vulnerabilities to Climate Change and Extreme Weather*, July 2013, (http://energy.gov/downloads/us-energy-sector-vulnerabilities-climate-change-and-extreme-weather).

10 *U. S. Department of Energy 2012 Strategic Sustainability Performance Plan, Appendix A*: *Climate Change Adaptation Plan*, 2012, (http://www1.eere.energy.gov/sustainability/pdfs/doe_sspp_2012.pdf).

11 *The Roadmap to Achieve Energy Delivery Systems Cybersecurity*, September 2011, (http://energy.gov/sites/prod/files/Energy%20Delivery%20Systems%20Cybersecurity%20Roadmap_finalweb.pdf).

Emergency Management Agency (FEMA).[12] These are but a few examples of DOE's commitment to increasing resilience of the energy sector.

This report includes a general resilience framework, including prototype metrics for analyzing, quantifying, and planning for energy infrastructure system resilience. Additionally, the report outlines the development of gas, oil, and grid specific metrics to provide a tangible example of how the general framework can be put into practical use.

[12] *Public Law No: 113-76, Explanatory Statement: Division D—Energy and Water Development and Related Agencies*, p. 64, January 2014, (http://docs.house.gov/billsthisweek/20140113/113-HR3547-JSOM-D-F.pdf).

2. WHAT IS RESILIENCE?

2.1 Literature Review

Many definitions of resilience exist, but few are standardized in their application to energy infrastructure. A literature review[13] on resilience and its application in energy and other critical infrastructures was conducted separately by the RAND Corporation, and should be referenced as a collaborative document.

2.2 Defining Resilience

The foundation of the work in this document is based on the definition of resilience offered through PPD-21, stating:

> "the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents."[14]

In refinements of this definition, the focus of resilience is on high-consequence, low-probability events (e.g., hurricanes, ice storms, malevolent attacks). This is not to limit importance to those areas only, but does recognize that many of the existing metrics in the areas of 'reliability' already include low-consequence, high-probability events (e.g., System Average Interruption Duration Index [SAIDI] or System Average Interruption Frequency Index [SAIFI] system-reliability metrics used in the electric power sector).

There is recognition too that in defining resilience, one should be able to measure it. As articulated by the National Academy of Sciences:

> "without some numerical basis for assessing resilience, it would be impossible to monitor changes or show that community resilience has improved. At present, no consistent basis for such measurement exists…"[15]

Finally, the recommendation is to use a risk-based approach to develop resilience metrics. This implies several key factors:

- Resilience is always defined with respect to a disruption or threat. For example, an electric infrastructure system may be resilient to hurricanes, but that says little about its resilience to ice storms, cyber attacks, or heat waves.
- Resilience metrics are defined to focus on the consequences of a system failure rather than the system failure itself. For example, it is possible that a portion of an energy infrastructure

[13] H. H. Willis and K. Loa, "Measuring the Resilience of Energy Distribution Systems," RAND Justice, Infrastructure and Environment, PR-1293-DOE, July, 2014.

[14] *Presidential Policy Directive 21, Critical Infrastructure Security and Resilience*, February 2013, (http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil).

[15] National Academy of Sciences, *Disaster Resilience: A National Imperative*, (National Academies Press, Washington DC, 2012).

system could be made more reliable (e.g., increased SAIDI or SAIFI system-reliability metrics), but the location of that reliability could have minimal advantage to social welfare during a major disruption.

- Resilience is always defined with respect to a specific system. That means that it is not useful to assign a value of resilience to a generic system, or one that has not been fully defined.

At a high level, a resilience metric must, at a minimum, consider the following key attributes:

- **Threat.** Definitions of likely disruption scenarios, with associated probabilities where appropriate.
- **Likelihood.** Probability that a disruption scenario may lead to decreased system performance or failure.
- **Consequence.** The impact of system failure given a disruption scenario.

We observe that reliability metrics do not possess these attributes—they are orthogonal in purpose and discrimination capability to resilience metrics. Additionally, we note that risk is informally (and often formally) defined as the arithmetic product of threat, likelihood, and consequence. As a result, a deep relationship exists between the notion of system risk and the concept of a resilience metric, such that changes in resilience will often (but not always) create changes in reliability.

As an example of these properties in a electric power sector context, consider the "threat" of disruptive weather to a particular electric utility. Likely disruption scenarios can be determined historically, defined in terms of the number and location of grid component failures. The likelihood (and extent) of each disruption scenario can be computed via systems analysis, given knowledge of available alternative generation and transmission/distribution resources and potential system states. Similarly, consequence for the disruption scenario can then be expressed in terms of the economic and/or public-health impacts. It should be clear that the definition of quantitative resilience metrics requires more data and stakeholder participation than reliability metrics, due primarily to the computation of consequences and the definition and scoping of low probability threat scenarios.

The precise definition of consequence is multi-attributed. In particular, consequence are quantified in terms of social impact (for example) loss of service, duration of loss, financial impacts of loss, recovery costs and resource requirements, and criticality of service where loss is incurred. Different operational decisions can yield different values of these quantities, and different stakeholders will emphasize different quantities. The decision-control architecture we introduce below facilitates comparison and analysis of the relationships between these various quantities.

System state is integral to definition and computation of energy resilience metrics, in particular because the extent of disruption due to a particular initiating event is highly dependent upon the system state at the time of failure. Further, because the certainty of our system state information is not perfect (e.g., due to measurement precision and limited quantities of sensors), resilience is

necessarily *stochastic* (i.e., the consequence of a particular disruption scenario is variable and uncertain). While the stochastic nature of resilience complicates the presentation of the resilience metrics, the information conveyed is more accurate and better represents the overall risk—and, as we discuss below, can therefore be the basis for control decisions to mitigate system risk.

Beyond system state, the final component in computing resilience metrics is the set of control decisions available to an operator and the values of those controls over time. For example, poor decisions about positioning of restoration crews in advance of a hurricane can significantly increase restoration time. In contrast, optimal decision-making in this regard may reduce restoration time to a few hours, minimizing the impact of the hurricane. Due to the relationship between operational decisions and consequence, it is impossible to separate the definition of resilience metrics from the specifics of the operational environment in which they are placed. This observation drives our discussion below of advanced decision-control architectures for resilience.

Above, we have outlined the key features of a resilience metric and the implications of those features for operational system control. Fundamentally, resilience metrics and the decision-control system in which they are placed are interdependent, such that static computation of a resilience metric cannot proceed independent of some representation of the control system. This has immediate implications for not only resilient operations, but also—and perhaps more importantly—investment planning.

Presently, grid operators possess no ability to quantify resilience during real-time operating conditions. Instead, available *reliability* metrics only indicate whether the necessary criteria are met for reliable operation. During electric-grid contingency situations that violate mandatory reliability criteria, such as the loss of multiple transmission lines or generators, operators have no means to quantify the risk to which the system is exposed. Without such guidance, power-grid control systems can neither effectively anticipate nor respond to contingency situations.

Figure 2. Critical Infrastructure Planning and Operations Timelines: The top-most vector represents operator knowledge and actions. The bottom-most vector represents various aspects of system design and operation.

It is useful to consider recasting system-specific operational régimes into this more general context. In **Figure 2**, we show the planning and operations timelines for a general critical infrastructure. A key point of this figure is to illustrate that post-event operational terms, including mitigation, response, recovery, and reinvestment play an integral and central role. In contrast, these terms have not made their way to the forefront of energy systems operations.

3. RESILIENCE ANALYSIS PROCESS

The Resilience Analysis Process (RAP) can be used to assess baseline resilience and evaluate resilience improvements. In simple terms, it explains how to 'use' a resilience metric. The process is designed to lead decision makers from high-level goals to a defensible, risk-based decision. **Figure 3** illustrates the RAP steps.

The first six steps of the RAP give decision makers and stakeholders a method for assessing the baseline performance of a system with respect to resilience. When all seven steps are followed, the focus of the RAP expands to identifying the improvements that will increase resilience. These improvements could be identified by analyzing or by optimizing the characteristics of these proposals to identify the best improvement strategies.

The RAP steps are depicted as a circle due to the iterative nature of resilience analysis. Francis and Bekera (2014) maintain that "vulnerability analysis at regular intervals is a key to recognizing disruptive events in advance and continuously self-evaluating and learning from incidents."[16] Periodic re-evaluation of system resilience is important for

- validating resilience analysis methodology,
- validating models against actual incident data, and
- updating resilience assessments with current technology methods and improved threat characterization.

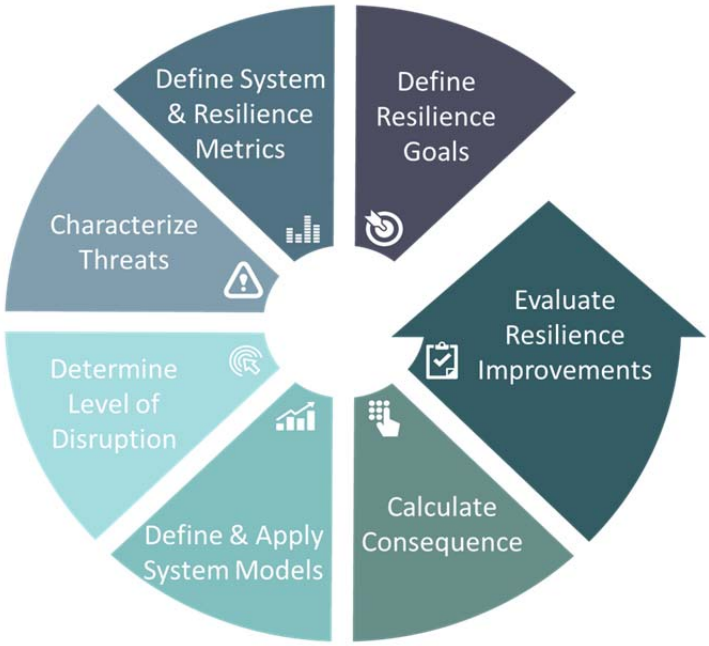

Figure 3. The Resilience Analysis Process: The steps of which progress counterclockwise from "Define Resilience Goals."

[16] R. Francis and B. Bekera, "A metric and frameworks for resilience analysis of engineered and infrastructure systems," *Reliability Engineering & System Safety*, **121**, 90–103, (January 2014), (http://www.sciencedirect.com/science/article/pii/S0951832013002147).

All relevant stakeholders should engage in the RAP. While there may be a single decision maker or small set of decision makers with a question of concern, engaging a broad community in each RAP step helps ensure a more complete analysis. Each step shown in **Figure 3** is described in the subsections below.

3.1 Define Resilience Goals

Before determining the scope of the system relevant for analyzing and selecting appropriate metrics, it is essential to define high-level resilience goals. The goal set during this first RAP step lays the foundation for all following steps. For example, discussion during this phase should determine whether assessing resilience is the main goal, or if evaluating possible system improvements is a central objective. If evaluating improvements is within the scope of the analysis, a decision should be made about the kinds of changes to be considered and the types of questions the analysis should address.

During this stage, key stakeholders and any possible conflicting goals are identified. Some examples of high-level goal language appropriate at this step of the process are:

- Improving a regional electric grid's resilience to natural disasters
- Deciding how to allocate a pipeline's capital investment and maintenance budget
- Ensuring availability of power to medical or transportation systems during disasters

3.2 Define System and Resilience Metrics

System and resilience-metrics definitions determine the analysis' scope. This could include identifying a system's geographic boundaries, relevant time periods, and/or relevant components. Because some consequence measures (e.g., macroeconomic impacts) require a relatively broad system definition, this is also a good time for stakeholders to discuss the types of consequences about which they are most concerned.

Determining the appropriate level of fidelity for the analysis should be driven at least in part by the high-level goals set during the previous step, although data availability can often drive these decisions as well. Metrics selected should be specific enough to enable decision-making, whether for operational or planning purposes.

3.3 Characterize Threats

Threat characterization is critical to understanding how capable the system must be to absorb and adapt to different types of attacks or natural events. When performing an analysis to evaluate resilience against multiple hazards, information about (1) the likelihood of each possible threat scenario and (2) the capabilities or strength of the threat is extremely important. In risk analysis, threat and consequence are used to understand which system vulnerabilities are most important to address to reduce the consequences associated with the threat.

3.4 Determine Level of Disruption

Once an understanding of the relevant threats has been solidified, the attributes of each threat are used to determine the amount of damage to the system (infrastructure, equipment, etc.) that is

likely to result from that set of threats. Models like FEMA's HAZUS can be useful at this stage—HAZUS contains models for estimating potential losses from earthquakes, floods, and hurricanes.[17] All of these types of models can be supplemented by noncommercial analysis such as optimization packages. This is the RAP step where expectations about structural damage or other system impacts that could affect performance are defined.

3.5 Define and Apply System Models

The damage states outlined in the previous RAP step can then be used as input to system models—tying damage to system output levels. For example, anticipated physical damage (or a range of damage outcomes when incorporating uncertainty) to an electric grid from an earthquake can be used as input to a system model that ties those outages due to damage to load not served within the system over time. Multiple system models may be required to capture all of the relevant aspects of the complete system. Furthermore, dependencies may exist between models. For example, a repair and cost model may be used to determine a repair schedule for components of an infrastructure. The schedule determined by these models may inform systems models to assess how the systems perform during the restoration period.

3.6 Calculate Consequence

When evaluating resilience, direct impacts to system output as a result of damage are only part of the story. Most energy systems provide energy not just for the sake of the generating or distributing it—but for some larger social purpose (e.g., transportation, health care, manufacturing, economic gain). During this step, outputs from system models are converted to the resilience metrics that were defined during the second RAP step. When uncertainty is included in the RAP, probability distributions will characterize the resilience-metric values. Depending on which stakeholders are performing the analysis, there could be significant differences about which kinds of consequences are prioritized. Methods such as the Analytic Hierarchy Process (AHP) can be used to support expert determinations of the relative importance of these measures.

3.7 Evaluate Resilience Improvements

Unless the RAP is being undertaken purely for assessment purposes, it is likely that some decision(s) must be made about how to modify operations or plan investments to improve resilience. After completing a baseline RAP through the preceding steps, it is possible and desirable to populate the metrics for a system configuration that is in some way different from the baseline in order to compare which configuration would provide better resilience. This could be

- a physical change (e.g., adding a redundant power line);
- a policy change (e.g., allowing the use of stored gas reserves during a disruption); or
- a procedural change (e.g., turning off equipment in advance of a storm).

Because the metric(s) and process for populating them are complete by this RAP step, all that is required is to repeat that calculation for the alternate system description.

[17] http://www.fema.gov/hazus.

Page intentionally left blank.

4. RESILIENCE METRICS

Standardizing a framework for developing metrics and a process for resilience analysis to use across different energy infrastructures allows for consistent decision making and sharing of lessons learned. Each infrastructure can then go even further and use a common set of metrics that are useful to most stakeholders within that domain. However, the RAP is not intended to be so prescriptive that it inhibits applying the general principles to unique problems. The RAP was intentionally designed to be flexible enough to allow for new or different domain-specific metrics while retaining the general principles that allow better assessment of resilience and evaluation of resilience improvements. This flexibility also applies to local- versus national-level problems.

4.1 Metrics: Measures of Resilience

The resilience metric framework that most effectively achieves these goals is a probabilistic assessment of consequence to extreme events, described in **Figure 4** as a probability distribution of consequence. The units of consequence are not specified in the framework, but must be defined during the resilience improvement process, resulting in a resilience metric. This probability distribution effectively accounts for the uncertainty associated with a limited understanding of future performance of the systems under threat.

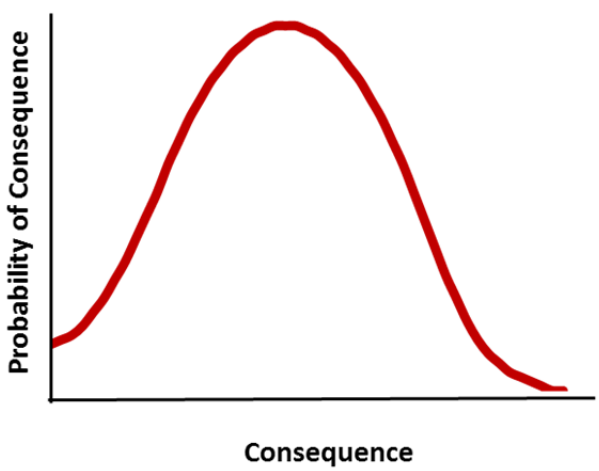

Figure 4. The resilience framework is a probability distribution of consequence to an extreme event(s). A resilience metric is a probability distribution of specific consequence to specific extreme event(s).

4.1.1 The Metric Is in Terms of Threat

Each resilience metric is presented in terms of a particular threat or set of threats. This tells the person using the metric what types of extreme events are being considered in the estimation of resilience (e.g. “resilient to what?”). More than one threat could also be represented by a single distribution. Only by being explicit about the threats being considered can the resilience metric quantify the uncertainty surrounding consequences of those threats.

4.1.2 The Metric Is Based on Performance

Although the resilience metrics will be used to assess the consequences of a threat, they must also assess how well the energy system fulfills its intended goals. This means that the resilience

metrics are directly related to system performance. Performance-based metrics are not directly based on system characteristics, but instead measure how well the system delivers on its intended purpose during and after the threat. For example, barrels of oil demanded but not delivered would be a measurement of how well the oil supply system performs after an earthquake. This is in contrast to attribute-based metrics (e.g., gallons of oil reserves available), which are complementary and useful, but not a sufficiently valid predictor of resilience. For instance, the number of spare transformers an electric utility has on hand improves resilience, but by how much? Without a measure of performance that is dependent on the number of additional spare transformers, the utility cannot make a cost-benefit decision on the number of spares to keep on hand. While often more difficult to populate, performance-based resilience metrics deliver the necessary information to make this type of decision.

4.1.3 The Metric Measures Consequence

Estimating or measuring the raw performance of energy infrastructures may not deliver the necessary information to quantify a system's resilience in the eyes of certain parties. For example, a municipality may care about how many people are able to continue working after an extreme event. In this example, the number of citizens able to work is a unit of consequence, and it is highly dependent on energy infrastructure performance. This is why the resilience metric framework calls for a transformation from indicators of system performance to measures of consequence. Consequence relates most directly to the fulfillment of social needs in these scenarios. It is likely that some users of this framework may choose to consider the system performance as the representative consequence, which is a reasonable simplification for some circumstances.

4.1.4 The Metric Accounts for Uncertainty

Understanding an infrastructure's resilience to a certain class of threats often means that most, if not all, possible instances of the threat have never been experienced by that infrastructure. To develop an estimate of performance and consequence in the event of these unexperienced threats, modeling the system will be necessary. These models will inherently have several uncertain parameters or relationships. By quantifying this uncertainty and propagating it through the models, a probability can be associated with each consequence estimate, thereby developing the probability distribution that is the resilience metric.

4.1.5 The Metric Effectively Captures Resilience

Given the working definition of resilience in PPD 21, it is necessary to show how the metric (a consequence probability distribution) measures a system's ability to prepare, withstand, adapt, and recover from potentially high-consequence, low-probability events. These four attributes and associated properties of resilient infrastructures, as described in **Figure 5**, are incorporated into the resilience metric by estimating their effect on the system performance, and ultimately their effect on consequence. For example, a system that recovers more quickly from disruption will have improved performance to one that recovers slowly. The improved performance is what is reflected in the metric, not the attribute that lead to that performance. Quantifying the impact of a system's attributes on performance will create a link between performance-based and attribute-based resilience metrics, and is therefore an important area for further research.

Capacities	Prepare	Withstand	Adapt	Recover
Example Infrastructure Attributes	Advance warning	Robustness	Rerouting	Mutual Aid Agreements
	Prepositioning	Redundancy	Substitution	Situational Awareness
	Stockpiling	Storage	Rationing	Resource Availability
		Separation	Reorganization	

Figure 5. Infrastructure system attributes that affect the resilience metric.

4.1.6 The Metric Is Not a Value Judgment

The resilience metric as stated here is not intended to be a declaration of requirements for any party. It provides a framework for assessing and improving resilience—as using speed measurements as a proxy for danger is a framework for assessing driver safety. This metric—a probability distribution of consequence—is a proxy for resilience, and it can be used to assess an infrastructure's resilience. It makes no supposition of what the metric's value should be for certain goals (e.g., it does not require the speed limit to have a 55 miles-per-hour threshold).

4.1.7 Multiple Metrics Are Often Necessary

Resilience metrics should be germane to the decision for which they are being used, and as a result, multiple metrics are often necessary. This could be the case when the decision is based on consequences that are either nontranslatable or when the translation is politically sensitive (e.g., loss of human life versus economic productivity). Many decision makers are already comfortable using multiple metrics with seemingly incongruous units. Often two or three metrics will be presented along a Pareto frontier, allowing for initial elimination of inferior alternative sets[18]. However, as the set of metrics becomes large—larger than three, for example—current methods struggle to present a completely objective set of comparable alternatives. This hampers the decision maker's ability to remain objective. Visualization of the multiple-objective consequence space with inclusion of probabilistic estimates of uncertainty represents a strong need for additional research.

Visualization of the multiple-objective consequence space with inclusion of probabilistic estimates of uncertainty represents a strong need for additional research.

4.1.8 System-Level Models in Resilience Metric Computation

A critical aspect of resilience analysis relates to the role of system-level models in computing resilience-metric values. We define a system-level model as a model that captures a system's fundamental operation—subject to perturbations or shocks induced by external events (e.g., hurricanes and earthquakes) or human-caused events (e.g., terrorism). This includes both

[18] V. Pareto, *Manuale di economia politica con una introduzione alla scienza sociale*. (Societa Editrice Libraria, Milano, 1906), English translation by A.M. Kelley, New York, NY, 1971.

delivery-of-service aspects (e.g., routing power in an electricity distribution grid) and recovery aspects (e.g., allocating crews to repair damaged equipment).

Because they must additionally account for operational behavior in regimes outside the standard-operations scope, the system models required for resilience analysis are different from those used in standard reliability operations. For example, while standard economic dispatch models used in power-grid operations account for the loss of a single system component, enforcing N-1 reliability[19], they do not capture decision-making processes under more extreme failure conditions. However, because standard operational models do capture key system behaviors (e.g., physical laws, in the case of the electricity grid), they are likely to serve as the basis for system models to enable resilience analysis. In contrast, system-recovery models are not standardized for most energy infrastructures, if they exist at all.

Across disparate energy infrastructures, the specific nature of system models for resilience analysis can be and are expected to vary widely. Depending on the fidelity required by a particular analysis, system-model complexity can vary from simple spreadsheet models, to more complex optimization models. Sometimes, more complex models can facilitate improved results, specifically reductions in resilience-metric values. For example, consider a simulation-based system-recovery model. Typically, these models encode specific operational decision sequences (e.g., prioritizating crews to repair specific system components). In contrast, optimization models may treat decision sequences as variables, such that the solution of those models can yield improved decision sequences relative to a simulation model. The net result can be a reduction in consequence, and therefore a reduction in the value of a resilience metric. This example highlights the critical linkage between the nature of a system model and the quantification of resilience—*system-level decisions impact the computed value of system resilience*.

Different resilience-analysis use cases can be accomplished with system models of varying fidelity. In planning contexts, coarse-grained models may be sufficient, as precise quantification of resilience is often not necessary. However, model fidelity must be significantly increased for both operations-oriented studies and operational use.

In conclusion, system models of post-event operation and recovery are foundational in resilience analysis and reside at the center of our proposed resilience framework. System models are required to compute resilience metrics, and the specific nature of the system models influences the specific values of the resulting metrics.

4.2 Resilience Analysis Use Cases

Metrics are necessary to make informed decisions when alternatives exist, to create goals, and to assess improvement toward those goals. Today, there is no agreed upon method to quantify the

[19] "System Performance Following Loss of a Single BES Element," NERC Standard TPL-002-0b, September 15, 2011.

resilience of a system or community.[20] This lack of resilience metrics makes creating resilience goals and driving systems toward those goals extremely impractical.

The resilience metrics put forth in this report not only allow quantification, but they enable system operators to increase control and optimization to increase resilience. New system operating paradigms can be established where resilience is employed to provide feedback in real-time to help minimize the impact of a catastrophic event on communities. Similarly, in the planning context, these metrics enable comparison between competing alternative investments based on resilience improvements, and design of optimal portfolios that maximize resilience under a constrained budget.

The resilience use cases presented in the next sections are examples of the types of new analyses and decision-support capabilities enabled by creating resilience metrics. These new capabilities can be classified in two main areas-- planning and operations-- and can be applied to all resilience definition phases: preparation, withstanding, adapting, and recovery.

First, the electric use case demonstrates the use of resilience metrics to quantify the grid's baseline resilience when operating through a natural disaster. Electric system operation is then shifted from economics-based principles to resilience-based principles and results are compared to the baseline resilience. A grid planning example is also discussed, where hardening of infrastructure options are evaluated under the same natural-disaster threat. Lastly, an example of optimal budget allocation is presented, where resilience is maximized for the same investment level as in the competing hardening options.

Second, a petroleum use case demonstrates the resilience metrics in the planning context by comparing the baseline resilience of the current system with that of an improved system through infrastructure investments.

Third, a natural gas use case evaluates the impact that operating policies for a natural gas storage facility have on the resilience of a region as measured by the impact that a natural disaster has in its economy.

All of the use cases are illustrative of the type of new capabilities enabled by the proposed resilience metric, highlighting the value of these new metrics. They also expose R&D gaps that should be addressed in the area of resilience.

[20] Committee on Increasing National Resilience to Hazards and Disasters and the Committee on Science, Engineering, and Public Policy, *Disaster Resilience: A National Imperative*, (National Academy of Sciences, Washington DC, 2012), 244 pp.

Page intentionally left blank.

5. A FRAMEWORK FOR DEVELOPING RESILIENCE METRICS

A resilience metric framework is defined as "the probability of consequence X given threat Y". The framework does not specify the specific threat or consequence and can be applied broadly. The red curve in **Figure 6**, below, is a notional representation of a resilience metric, and because it does not specify the specific threat or consequence, we refer to it as a metric framework.

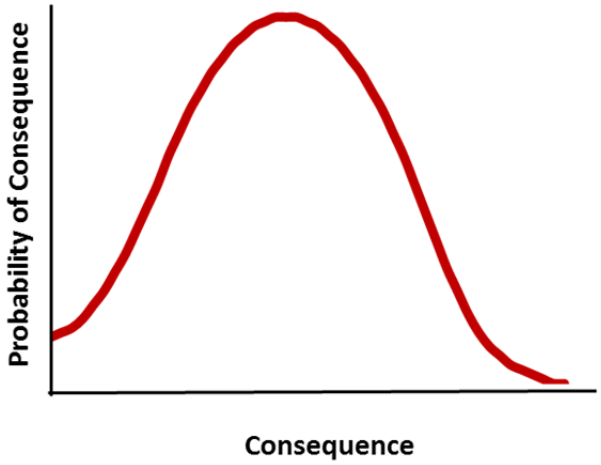

Figure 6. Resilience Metric Framework.

Resilience metrics should be able to accomplish several key elements. The following requirements provide a means for developing a resilience metric framework. They must:

- **Be useful.** Metrics developed under this framework must be useful for decision making (by humans, computational analysis, or both). Decisions of interest include system planning decisions, real-time operations decisions, and policy decisions.
- **Provide a mechanism for comparison.** Applying the same metric to different systems should result in valuable information. Furthermore, the same metric must be able to differentiate between the resilience of a system that has not been hardened and one that has (either through infrastructure or operations enhancements).
- **Be useable in operations and planning contexts.** The same metric should be able assist decisions in both operating conditions (such as preconfiguring a system before a hurricane) and planning (such as burying electrical conductors).
- **Exhibit extensibility.** The metrics selected must be scalable in time and geography. The metrics should remain valid as technology progresses and more complex analytic methods become feasible.
- **Be quantitative.** The framework must allow the development of metrics that can be used both qualitatively and quantitatively.
- **Reflect uncertainty.** It's critical that metrics are populated using methods that will quantify the uncertainty of the result. Specifically, decisions being made based on a resilience metric value must be well informed by the certainty of that value.

- **Support a risk based approach.** The metrics should reflect a specific hazard or set of hazards, the system vulnerability, and potential consequences to people (beyond the immediate system effects).
- **Consider recovery time.** Resilience metrics should reflect the amount of time that an outage occurs, either directly or indirectly.

It is important to highlight a distinction between unpopulated metrics and populated metrics. One example of an unpopulated metric is *miles per hour*. This metric might be useful for measuring the speed of a vehicle, an airplane, a racehorse, or even an asteroid. Populating this metric might be simple (using a speedometer, in the case of a vehicle), or quite complex (advanced analysis, in the case of an asteroid) and the process of populating them need not be the same. The analytics associated with populating the metrics are critical, but strictly speaking, they are not part of the metric. Improvements in technology and analytics may allow improved certainty of a populated metric, but have no effect on the metric itself.

In contrast to a metric that has been populated as an assessment of a system, goals are policy decisions that are used to populate metrics. For example, 55 miles per hour is a goal (or limit), which may be compared with the populated value of speed found on your car's dashboard. These distinctions are important as this report discusses separately topics on (1) metrics, (2) analytic methods for populating metrics, and (3) resilience goals. The ability to populate resilience metrics and the decisions made based on them (e.g., policies, goals, and other actions to minimize the consequences of a hazard) are all extremely important and will be discussed in a later section of this report.

6. POPULATING METRICS

Resilience metrics take the form of a probability density function (PDF) of the consequence of interest. **Figure 7** and **Figure 8** provide examples of resilience metrics when economic impact and safety are the areas of interest. There may be many approaches to obtaining these PDFs with diverse levels of complexity. For instance, a complex approach might take into account many sources of uncertainty which may include uncertainty models that produce conditional probabilities for each realization, while a less complex approach might consider fewer sources of uncertainty and assume their realizations are independent. An even less complex analysis could assume there is a high degree of certainty in all of the analysis inputs and calculate a deterministic value while performing a sensitivity analysis to understand how small changes in those input values change the results. On the other hand, the presentation of the results might also be reduced by communicating features of the PDF such as mean value or Value at Risk (VaR).

Regardless of the complexity of the analysis method employed or the form in which results are communicated, resilience metrics can still be used to help operate and plan systems in a way that improves their resilience to high-consequence, low-probability events.

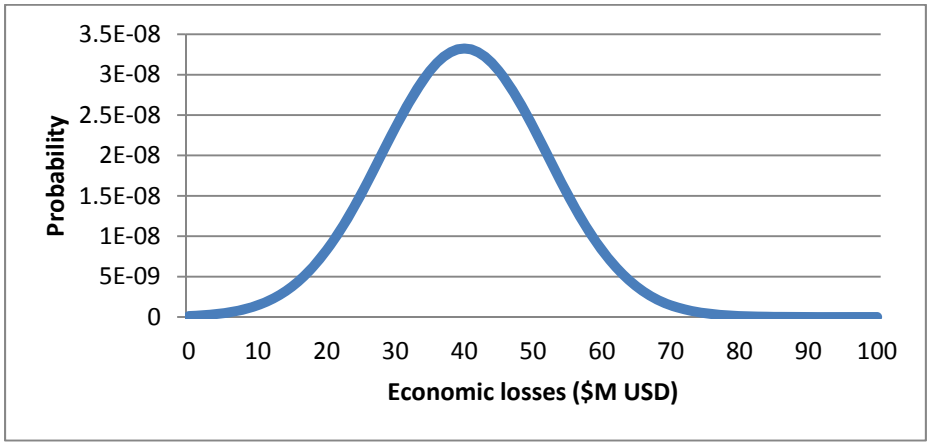

Figure 7. Example of a PDF describing resilience in terms of economic impact ($M USD).

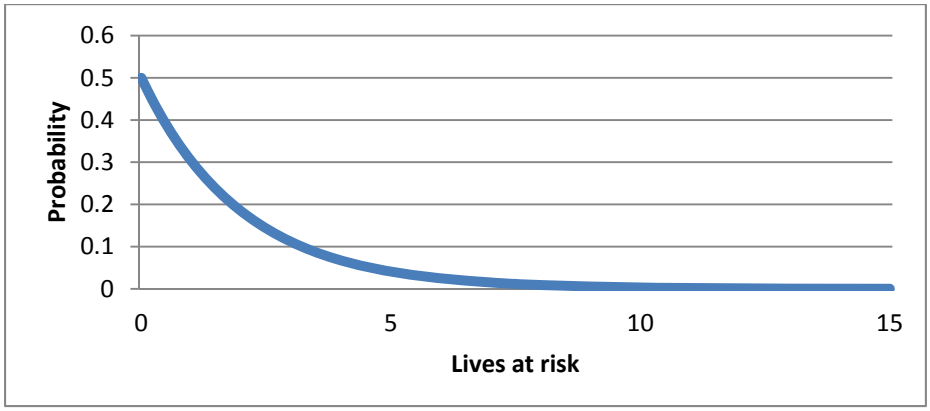

Figure 8. Example of a PDF describing resilience in terms of lives at risk.

Page intentionally left blank.

7. SIMPLIFIED, ANECDOTAL USE CASE

This section describes how a hypothetical decision-making process could be augmented by using resilience metrics, and how those resilience metrics would be calculated and used.

7.1 Define Resilience Goals

In this anecdotal use case, the decision maker is a public electric utility that is wholly owned and operated by a municipality. This utility is making a planning decision related to improving their built infrastructure and is being asked by the mayor to explicitly account for the possibility of extreme events. The utility's stated goal, in this respect, is to deliver energy at reasonable cost, with minimal negative impact to public productivity. They have experienced one hurricane in the past that caused public productivity to drop to unacceptable levels in the eyes of the mayor's office, so they are primarily interested in improving resilience to future hurricanes.

*The utility's stated goal is to **deliver energy** at a **reasonable cost**, with minimal negative impact to **public productivity**, accounting for the possibility of extreme events such as hurricanes.*

7.2 Define System and Resilience Metrics

To define the system and metrics, the utility uses the stated goals and determines indicators of performance that will quantify how well these goals are being met. Ultimately, the decision of which resilience improvements in which to invest will be made based on an apples-to-apples comparison. So, the utility also must determine the units of consequence for the resilience metric. In this case, performance is measured by three indicators as described in Table 1:

1. the capacity unavailable to serve load,
2. the increase in cost of operation, and
3. the decrease in labor hours by the public over the course of the recovery period.

These performance indicators are dependent on grid behavior as well as the economic conditions for the utility and the community. Multiple systems will be assessed in order to calculate these indicators, such as the physical grid, the economics of grid operations, and the status of businesses and the labor force. The consequence will be measured by the overall economic impact to the community, which will be calculated using these performance indicators. This means that the resilience metric will be a probability distribution of economic impact. For this anecdotal use-case, the method of conversion of each performance indicator to economic impact is not addressed. However, it will be addressed in the detailed examples later.

Table 1. Mapping between goals, performance indicators, and their definitions

Goal	Performance Indicator	Definition
Deliver energy	Unavailable supply	Nominal supply—Actual supply
Reasonable cost	Increase in operation cost	Operating costs—Nominal costs
Public productivity	Decrease in labor hours	Nominal labor hours—Actual labor hours

7.3 Characterize Threats and Determine Level of Disruption

As mentioned, this utility has experienced one hurricane in the past, and the consequences were dire enough to spur the interest in resilience-enhancing electric grid investments. While they are most interested in potential hurricanes, they also recognize that other threats should be addressed for a more holistic understanding of resilience. In this case, the utility assembles data on how the hurricane disrupted the systems of concern, and how it impacted the performance indicators. They use this data to calibrate models of future hurricanes and how these systems may be disrupted.

7.4 Define and Apply System Models

The utility starts by calculating the performance indicators based on historic data from the past hurricane. By the definitions of the performance indicators in Table 1, the utility compares performance during the hurricane to nominal conditions, e.g., a statistically equivalent situation in which the threat was not present. They generate graphs of performance over time as presented in **Figure 9.**

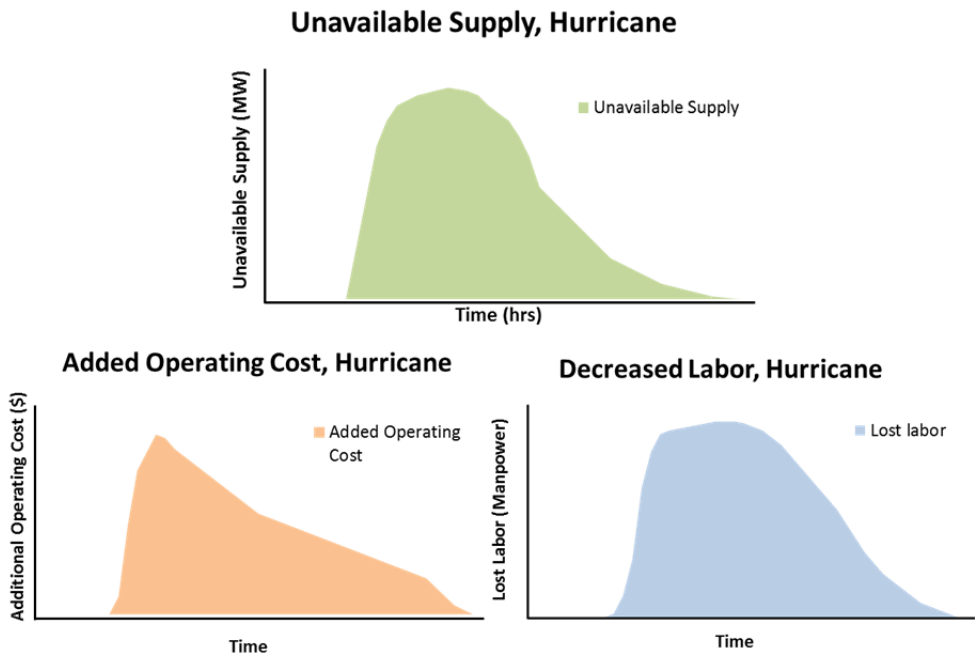

Figure 9. Performance indicators through time for one particular hurricane instance.

To the extent possible, the utility considers sources of uncertainty that would impact how their system might perform in the face of the next hurricane. They could be uncertain about the damage due to the hurricane, the intensity or path of the hurricane, the response of the public, or the resources available to them for repairs, to name a few factors. These uncertainties must be quantified to the best of their ability and then propagated through the models to understand the probability associated with the multiple estimates of performance. At this point, with multiple model runs indicating multiple possible futures and their associated probabilities, the utility is ready to calculate consequence.

7.5 Calculate Consequence

The utility understands and models how the performance indicators translate to consequence in their system. Using this transformation algorithm, they come up with time plots of economic impact for each realization of a possible future from their multiple sets of performance indicators, shown at left in **Figure 10**. Using the probability associated with each of these possible futures, the accumulated consequence, which is measured in economic impact in this case, is plotted as a probability distribution, as shown at right in **Figure 10**. The utility now has a probabilistic assessment of consequence to hurricanes—one which helps them understand the potential economic consequences. Based on this metric, they can decide if they need to take additional action, such as system hardening to reduce these potential consequences, either the expected consequences (mean) or the value at risk, (the tail).

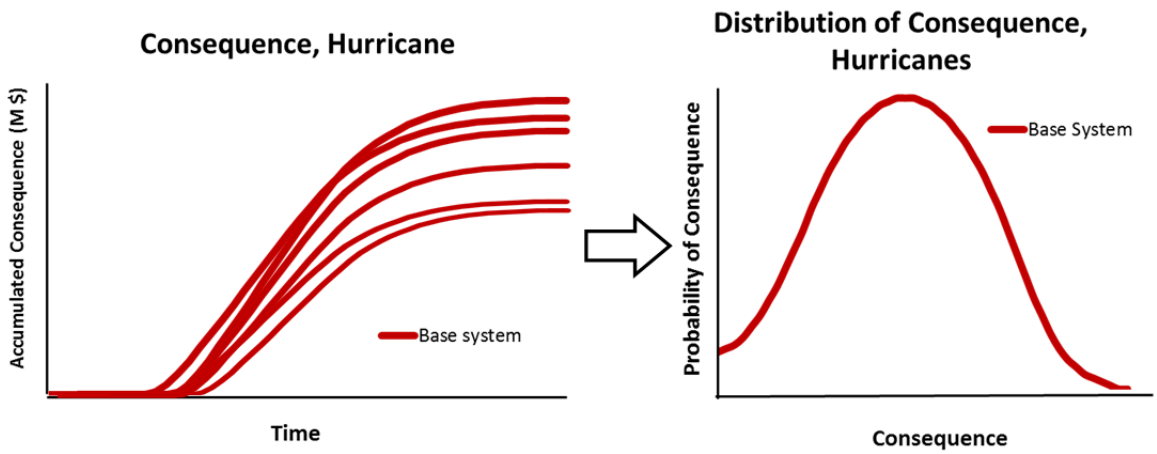

Figure 10. After transforming the performance indicators into consequence and accounting for uncertainty (left), the consequences for all realizations are plotted as a histogram of consequences (right). This histogram in converted into a probability density function, and is a resilience metric.

7.6 Evaluate Resilience Improvements

The metric presented in **Figure 10** is a measure of the existing system's resilience to hurricanes, but resilience could be improved if investments were made to modify the system. Each investment alternative will alter the system's characteristics to be more resilient, perhaps by recovering faster, adapting more quickly, or mitigating initial impacts, and each alternative will have an associated cost to the public utility. To analyze the alternative investment scenarios, the utility would identify feasible investments that might be expected to increase the resilience (such as burying overhead conductors, adding floodwalls, or oversizing various equipment). These alternative investment scenarios would use the same modeling techniques just presented, but would first modify the system model to reflect the new improvements. For each particular alternative, they would calculate a new ensemble of disruption estimates, performance indicators, and finally a new resilience metric. By comparing the metrics associated with investment alternatives and to the do-nothing scenario as illustrated in **Figure 11**, they would decide if and how they would invest to increase their resilience.

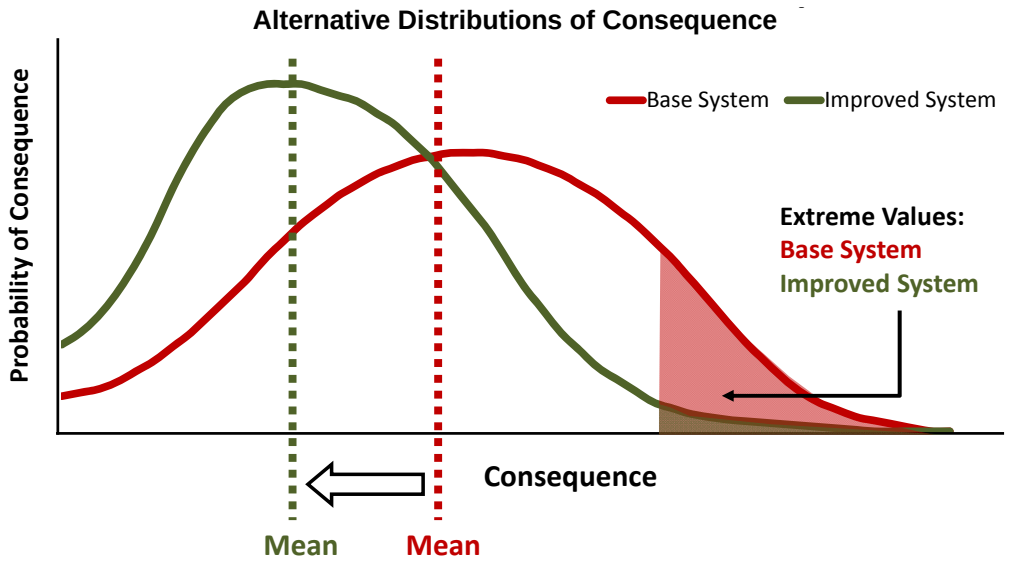

Figure 11. To analyze an investment alternative for resilience improvement, the resilience metric for the unimproved system is compared to that for the improvement alternative.

The statistical mean of each distribution indicates the expected consequence to future hurricanes, so perhaps the utility would consider investing in the alternative if the difference between the mean for the alternative and the mean for the base case outweighed the project's cost. In other words, the utility might expect an alternative to be worthwhile if it were expected to improve overall economic impact by more than the cost of that alternative. There are other ways to make decisions by comparing resilience metrics. The area of the distributions greater than a particular consequence value, VaR, could also be compared, as is pointed to by the extreme values arrow in **Figure 11**. The considered investment alternative in green decreases the probability of extreme consequences considerably, as indicated by comparing the two shaded areas under the curves.

7.7 Incorporating the Resilience Improvement Process

Using the resilience-metric framework, this hypothetical utility has determined an investment strategy that will not only improve resilience, but will do so cost effectively and will fit within their overall goal of delivering energy at a reasonable cost with maximum benefit to public productivity. It is imperative to include this resilience-improvement process within the overall planning process because the considered alternatives will also have an impact on day-to-day operations and associated metrics such as reliability and efficiency.

8. ADVANCED USE CASE SUMMARIES

The following section describes three specific cases in which the resilience metrics are analytically populated and used for decisions that are commonly found in industry. One use case for each of the three major energy sectors is assessed here: electric power, petroleum, and natural gas. Each use case follows the resilience-enhancing process presented in Section 7, populating a resilience metric as a critical part of this process. Table 2 outlines the threats, performance indicators, and measures of consequence that the use cases consider.

Table 2. A summary of the specific aspects of the resilience-enhancing process that each use case covers

Sector	Threat(s)	Performance Indicator(s)	Consequence(s)
Electric Power	Hurricane	MWh not served	Economic losses
Petroleum	New Madrid Earthquake	Barrels of fuel not consumed	Added cost of fuel
Natural Gas	San Andreas Earthquake	Mcf of gas not delivered	Economic impact

The three cases overlap in several dimensions. Each case focuses on one type of threat, one performance indicator, and one measure of consequence—although this is not a necessary condition of the resilience metric framework. The electric power use case presents the use of a resilience metric to improve planning as well as operations decisions. The petroleum case focuses on an investment decision about pipeline modifications, while the natural gas case highlights a policy decision about how to manage storage. The cases show that while the process of populating and using the resilience metric can be standardized, the decisions to which it can be applicable as well as the type of models and algorithms necessary are wide-ranging and flexible.

8.1 Electricity

The electricity use case is an illustrative example of resilience analysis for the electric grid, using the proposed resilience framework and associated metrics. The analysis is organized around a series of mini use cases—ranging from

- a baseline resilience computation for an existing system to
- a comparison of alternative investment portfolios to enhance resilience to
- an optimization of investments for enhanced resilience, given a fixed budget.

Threat scenarios are defined to include hurricanes and consequences are shown as “demand not served” in one case and “economic losses” in another. For more details, see Appendix B.

8.2 Petroleum

This use case demonstrates one way to use the resilience framework to identify potential options to increase resilience and measure the increase in resilience due to implementing these options. Specifically, calculate the increase in resilience gained by re-engineering two major transmission pipelines to decrease down time after a large scenario earthquake in the New Madrid Seismic Zone. For more details, see Appendix C.

8.3 Natural Gas

The use case presented exemplifies how resilience metrics can be applied to the natural gas infrastructure. Resilience is evaluated by calculating the overall impact on the economy ($) that natural gas delivery shortfalls would cause due to a natural disaster, a magnitude 7.8 earthquake at the San Andreas Fault near the Salton Sea. An engineering assessment for this type of earthquake was performed and results show that it would damage three important transportation corridors around the southern California area. For more details, see Appendix D.

9. WHAT CAN BE SIMPLIFIED TODAY

When modeling a system, there are two extremes in terms of complexity:

- a model of high fidelity (and complexity) that provides a comprehensive output and
- a model that is so simple that it can rarely be applied to real-world problems.

Although both of these models have their place, often, the most useful model lies somewhere in between the two ends of this spectrum. The same can be said about the models employed for populating resilience metrics. More data and more refined models will result in high-fidelity results, but many models that require less data or are less computationally intensive are also useful.

The appendices present several illustrative use cases where operating and planning computer models were employed to quantify resilience, perform control actions, choose between competing investment decisions, and form an optimal portfolio of investments in three different energy infrastructure types. These were developed to showcase a range of new capabilities that resilience metrics enable. And while their effectiveness depends on model accuracy, validity, and fidelity, these models are useful even when some of the analysis inputs and tools are simplified. A few simplifications that can be made using the metric-framework and analyses presented. They include:

- Use historical data from previous events to replace computer-based models
- Employ SME information to simplify threat characterization, reduce uncertainty sources or facilitate system analysis
- Employ resilience-based strategies compiled in “best-practice” manuals or “playbooks” instead of real-time tools
- Reduce the number of scenarios that describe threats or uncertainty
- Simplify resilience-metric presentation from PDFs to key values on the PDF (mean; percentile, Value at Risk)

9.1 How we can prepare for a more resilient framework

The process can be used to support many different decision types—both strategic and operational—as well as by different entities for different purposes. It could be used in support of federal, state, and local policy decisions. It also provides decision-makers with a methodology that is transparent, traceable, and defensible—ensuring choices can stand up to the intense scrutiny that accompanies the allocation of limited, publicly funded resources. Within the private sector, this framework could provide support for decisions about where to deploy protective measures, create redundancies, or develop response plans.

Another benefit of the process is with regard to complexity, insofar as it presents a method to evaluate resilience, but can be useful for varying degrees of analytic and modeling maturity. By design, the process is flexible. Not every step may be necessary for every system being analyzed. In addition, the tools and techniques are not intended to be exhaustive. As appropriate, analysts

should apply other tools based on the specific situation, the cost of implementing the tool, and the quality and type of analysis needed. The process also provides a roadmap for the research community to begin incorporating more advanced concepts as R&D matures.

10. RESEARCH & DEVELOPMENT

10.1 Furthering Resilience

As described in Section 3, the RAP identifies seven key steps for assessing and improving energy infrastructure resilience. To provide a more robust capability over time, specific steps in the RAP will require additional R&D. As shown previously in this document, achieving improvements to the resilience of a given infrastructure requires varying levels of analytic representation depending on the threat and the infrastructure in question. RAP steps 2 through 7 can be analyzed with SME input and simple analytic representations. However, due to the complexity of these systems and the inherent uncertainty associated with the infrastructure threat, as well as the uncertainty of infrastructure disruption, more sophisticated analytic approaches are normally required. The desire to understand these complex systems in more detail—so that more robust, complete assessments can be obtained—is the basis for material improvement of the US energy infrastructure. A major part of this R&D effort must be ensuring this framework is practicable, useful, and ultimately adopted—or at a minimum, strengthens and informs resilience analysis and decision processes.

The first step in assessing an energy infrastructure's resilience is defining appropriate resilience metrics and calculating a quantitative measure for those parameters. Section 10.2 focuses on the R&D aspects of this important step. Once the energy system and its associated resilience metrics have been identified, crucial information is still required to populate them (i.e. evaluating the system's overall resilience). Additional R&D is required in these areas as well.

The primary elements of an R&D program in energy resilience should include further advancements in the following areas.

10.1.1 Characterize Threats

Energy infrastructures are vulnerable to multiple threats, including both natural and man-made. More research is required to create better ways of evaluating these threats specifically focusing on the uncertainty associated with them. Also, not all threats are equally likely, and not all affect energy infrastructure the same way. For instance, a wind storm can cause considerable damage to the electric infrastructure but might have no effect on petroleum or natural gas infrastructure which is largely underground. For these reasons, it is important to include threat characterization as part of the R&D program.

10.1.2 Determine Level of Disruption

Once the threat to an energy infrastructure has been sufficiently characterized, the estimated level of disruption to the system must be determined. Though some capabilities currently exist for some structures (e.g., HAZUS), additional capabilities are required to estimate disruption levels to specific infrastructures. For example, as a natural disaster hits a region, uncertainty can exist about the extent of the damage it will cause to energy infrastructure system components. Analytic models must be developed and historical data collected to create capabilities that can accurately predict system conditions after a high-consequence, low-probability event occurs.

10.1.3 Define and Apply System Models

To accurately determine a system's resilience, good system-level models are required—preferably with as much fidelity as possible. Further R&D is required to continue to improve the ability to describe these complex systems with more precision so as to improve the quality and fidelity of the assessments for decision makers. Specific attention should be given to system models that allow the removal of presumed portions of the system. Given an assessment of the disruption level caused by a particular threat, a robust system model must be able to effectively remove impacted portions of the system. With these components removed, the system model must be able to function reliably to determine the overall impact to the system performance. In addition, better system characterization after impact will allow for procedures and practices that will help energy infrastructure better withstand those event types. Additional R&D is required to improve these capabilities.

10.1.4 Calculate Consequence

As stated earlier, most energy systems are providing power for some larger purpose (transportation, health care, manufacturing, economic gain). Further research is required to determine appropriate analytic techniques to incorporate the uncertainty inherent in this RAP step. Various mathematical techniques are available, but further investigation is required to determine the appropriate methods and their effectiveness. A comprehensive research program will benefit from including research in this area.

10.1.5 Evaluate Resilience Improvements

This step focuses on how to modify operational decisions or plan investments to improve resilience. Further R&D is required to effectively quantify uncertainty for these evaluation processes. Computational methods should be developed to provide more robust methods for creating and evaluating proposed resilience improvements.

It is important to note that the entire process to assess and improve an energy infrastructure's resilience requires human involvement. From an operator executing procedures to an executive making financial-investment decisions, the behavior of humans is an important part of determining the resilience of an energy system. For this reason, further R&D is required to incorporate the human-in-the-loop aspects of each of the RAP steps.

In summary, a robust, sophisticated energy infrastructure resilience R&D program should include investment in

- more advanced, quantitative analytic methods;
- the ability to incorporate uncertainty into all aspects of the resilience-problem domain; and
- the incorporation of human behavior into the assessment process.

Such a research program should be integrally connected with a strong stakeholder engagement activity.

11. MOVING FORWARD

11.1 Identifying Specific Metrics that Could Be Generally Applied Across Many Different Systems

Resilience is critical to stakeholders at many different levels and throughout many different infrastructures. Standardizing a framework for metric development and a process for resilience analysis for use across different infrastructures allows for consistent decision-making and sharing of lessons learned. Each infrastructure can then go even further and use a common set of metrics that are useful to most stakeholders within that domain.

However, the RAP is not intended to be so prescriptive that it inhibits application of the general principles to unique problems. This process was intentionally designed to be flexible enough to allow for the use of new or different domain-specific metrics while retaining the general principles that allow for better assessment of resilience and evaluation of resilience improvements. This flexibility also applies to local versus national level problems.

Selecting the appropriate metric for a given analysis should be driven by the stakeholder goals. Metrics in this context should be selected for their ability to enable resilience decisions or assessments. However, data availability is often a constraint, so comparison of data that is already being collected about the system in question and possible metrics under consideration may help focus early efforts when funding and time are limited. Data-collection needs can then be documented and planned for in later cycles.

While the selection of specific consequence-based metrics should be based on the goals identified by stakeholders in the first step of the RAP, common system output measures in each infrastructure are likely to be useful regardless of the consequence metric. One example of such a performance indicator for electric power would be “load not served” over the course of the disruption. Tying these performance indicators to relevant measures of consequence, like economic loss or hospital beds unavailable due to power loss, is a challenge that would greatly benefit from further R&D.

One advantage of pursuing these research activities at a national level is that it will help ensure consistency of methodology and enable discussions between regions or localities. A library of suggested performance indicators and recommended methods for translating those system outputs to some common consequence measures is a worthwhile and necessary national R&D pursuit.

11.2 Stakeholder and Buy-In Process

An important first step in using resilience metrics to improve nationwide energy resilience is to vet them and insert them into the energy industry’s lexicon. This will require a concerted engagement process that will consider how the metrics will ultimately be used and by whom. The resilience-metric framework fulfills the needs of multiple parties in this regard, so multiple conversation threads will be necessary to capture the full breadth of potential uses. During this process, it may be important to highlight that the resilience metrics themselves do not necessitate a change

in business practices, but instead augment current practices by fitting well into existing cost/benefit and risk-management frameworks.

As the resilience metrics are vetted and ultimately become part of the energy-industry lexicon, focus will be placed on how the resilience analysis and improvement process will be implemented. The goal of this implementation will be defined by multiple stakeholders in both the public and private sectors, but should ultimately focus on improving energy infrastructure resilience with respect to the consequences chosen by those stakeholders and in consideration of threats chosen by those stakeholders. The benefit of the resilience analysis and improvement process as it is presented herein is that it provides a guide for measuring resilience, but allows flexibility for different communities or groups to determine the consequences and threats relevant to them.

A key goal of the resilience-improvement process is to facilitate communication—especially among different types of decision makers and across interdependent energy infrastructures. This will not be possible unless decisions are made as to how to populate metrics. Units of consequence, approaches to uncertainty, and consideration of threats will need to be decided upon in an inclusive manner for multiple types of resilience-improvement decisions as well as multiple regions.

There will be situations where the units of consequence must be standardized so that resilience can be compared across multiple parties, such as in a regulatory or common-investment decision process. For instance, there may be a situation where a government entity secures funding for an investment in energy resilience, but wants to do so equitably and with greatest impact to community resilience. To compare alternatives for this investment, the units of consequence will need to be the same across all analyses and should also be normailzed. These units should be informative for the government entity, but also be feasible for potential bidders on these projects to calculate. In a different scenario, a regulation may be placed that requires a certain level of resilience, but the reporting requirements to prove compliance should not place undue burden on the reporting parties. In this case, the units of consequence will be related to what the regulated parties can feasibly measure and calculate. By considering the capabilities and data collection performed by the parties affected by the resilience-improvement process, it will complement, not replace existing risk frameworks for asset protection.

Similar to the considerations for the units of consequence, consideration of threats will be important to gain legitimacy for the resilience improvement process. Just as units of consequence may need to be standardized for different types of decisions, the threats considered will likely vary regionally. The considered threats will take into account the probability of that threat in the region as well as the potential consequence. For instance, analysis of resilience in the Midwest will likely consider tornados, while a similar analysis on the West Coast may not need to consider this threat. A regional process for which threats to consider will need to be developed that includes objective information from multiple parties about the likelihood and potential consequence of multiple threat types.

Finally, standardization will be important for incorporating uncertainty in the resilience metrics. This process may not involve a complete set of data, so expert judgment can be used where data does not exist. Some categories of uncertainty will be more controversial to quantify than others. For instance, quantifying the uncertainty surrounding hurricane path and intensity may lead to little controversy, while uncertainty associated with the economic impact of loss of energy to one neighborhood versus another could create substantial controversy. It is important to understand which uncertainty categories are most critical to estimate accurately, and which are likely to cause controversy. The former criterion could be assessed using modeling and sensitivity analysis, while the latter may require polling. When a category fulfills both of these criteria, it will require careful attention and buy-in from multiple affected parties. These categories will most certainly require a public processes.

11.3 Development of Sector-Specific System Models

An integral aspect of the proposed resilience framework is the underlying system models, which capture the behavior of the infrastructure when subjected to a threat. These behavioral models are in contrast to standard, typically reliability-centric operations models. In particular, for a given infrastructure, it is necessary to specify recovery and restoration processes, in addition to processes invoked by operators when a disruptive event is unfolding.

Such models do not presently exist for the electricity, petroleum, and natural gas infrastructures. Thus, a key R&D challenge in the application of the proposed resilience framework consists of development of baseline system models. We expect the baseline models to initially involve extensions of existing reliability-oriented system models, which focus strictly on delivery sustainment—typically at minimal cost. Significant extensions of these base models will need to be developed, particularly in the areas of recovery and restoration modeling. The purpose of the new models is to facilitate resilience-metric computation, through a reasonable-fidelity model of the underling system and the consequences associated with loss of delivery.

Development of resilience system models should proceed in partnership with industry, to ensure their acceptance and ultimate adoption. The models are likely to vary in fidelity, depending on the specific type of resilience analysis (e.g., planning or operations) being considered. For purposes of long-term adoption, development should initially focus on low to moderate-fidelity system models, to ensure that key properties are captured and the range of target resilience analyses can be conducted. Then, as the initial models are better understood, fidelity can be incrementally improved.

11.4 Consequence Quantification

Another major component of the proposed resilience framework involves specifying a given infrastructure's loss-of-delivery consequences. Two key issues dominate the specification process: (1) identifying high-level consequences and (2) computing loss-of-delivery consequences.

Identifying high-level consequences involves defining one or more consequences critical to a particular resilience analysis. For example, a particular stakeholder may be interested in resilience as expressed through both "safety" and "economic" factors. However, significant research

and practical challenges are associated with precisely defining such factors. Further, because the precise definition is stakeholder-dependent, any industrial instantiation of the proposed resilience framework must proceed with an industrial or governmental partner. Initial pilot studies will gain insight into the difficulties in defining consequence along a number of informal dimensions, which will aid further adoption of the framework.

Computing loss-of-delivery consequences for a particular system involves either leveraging or developing a deep understanding of the dependent infrastructure served by the system under consideration. As with identifying high-level consequences, a significant portion of the process is social and involves interaction with domain experts. In this phase, the system must be analyzed from a loss-of-delivery consequence perspective, with appropriate parameterizations made for inclusion into resilience system models. Again, such analysis will necessarily proceed with an industrial or governmental partner, to promote adoption and to identify issues associated with subsequent applications of the proposed resilience framework.

11.5 Conclusions and Recommendations

A framework for energy resilience metrics has been created such that:

- Energy resilience metrics quantify the expected consequence due to events that have low probability but potentially high consequence. Consequences focus on social welfare, extending beyond system impacts.
- The resilience metrics are based on the performance of the system, as opposed to being attributes of that system.
- The resilience metrics incorporate the uncertainty associated with limited information about the system and the threat.
- Resilience metrics quantify performance given uncertainty, providing insights into risk management and cost/benefit processes for planning, operations, and policy building.

A resilience analysis process has been created that explains how to use resilience metrics. The process is flexible enough for use by different stakeholders and infrastructures. Stakeholder goals should drive the selection of metrics used for an analysis within the framework provided.

Continued research is essential:

- More research is needed to improve quantification of human/societal consequences based on reduced system performance in a disruption. Key areas for R&D investment include multi-category uncertainty quantification, modeling and simulation of disruption, recovery and repair, communication of risk, and adaptive system operation algorithms.
- Developing a library of suggested performance indicators and recommended methods for translating those system outputs to common consequence measures is a necessary national research and development pursuit.
- Data availability will be a challenge in the early stages of adopting these methods, so some effort is likely to be needed with respect to data collection and establishing associated best practices.

Outreach and collaboration is necessary to define the types of decisions that will use resilience metrics, as well as the metrics' units of consequence, selection of threats, and quantification of uncertainty.

A stakeholder group should be created for the refinement and standardization of metrics for electricity, petroleum, and natural gas for the validation of this resilience metric framework. Specific areas that should be addressed include:

- Differentiate reliability metrics from resilience metrics with input from state, federal and regional regulatory authorities and other stakeholders
- Determine federal, state, and local government roles
- Work toward stakeholder buy-in and coordination: federal and state regulators, utilities, asset owners, and other key stakeholders
- Conduct an expanded case study using data from a major utility (in coordination with that utility)

Page intentionally left blank.

APPENDIX A: RELIABILITY AND OTHER COMPLEMENTARY METRICS

This appendix has been added to describe the very mature set of reliability metrics that already exist and are routinely used as the basis for regulatory assessment and rate recovery. The mixing or confusing of reliability metrics with resilience metrics, therefore, should be avoided so to avoid substantial opposition from those using reliability metrics. This appendix highlights that reliability metrics are nearly always attributed to high-probability, low-consequence events. Even for metrics such as SAIDI and SAIFI, reported values of these statistics are scrubbed of data that result from severe storms or other high-consequence events so that public utility officials can use these statistics to compare normal operating reliability across several utilities. Another distinction between reliability metrics and the resilience metrics is that reliability metrics do not attribute cause to the metric (a load is de-energized but without regard to why or how), whereas resilience metrics do (e.g. a hurricane caused the load to be de-energized). We note that systems may be considered reliable without specifying what is threatening the system, but when discussing resilience, systems are always resilient to a particular threat or set of threats. Finally, reliability addresses the ability of a system to accomplish its objective; which says nothing about how the system response may affect the community or other social elements. Again, resilience bridges this gap by extending the system response to a social conclusion.

In a general context, reliability is the ability of a component, device, or system to perform its intended function. Related to power systems, reliability is assessed based on how well the system supplies electrical energy to its customers. There is a tradeoff between how reliable the power system is and the investment needed to achieve or maintain reliability levels. This is illustrated in Figure 12 where the change in incremental cost of reliability is depicted as the ratio of the change in reliability ΔR to the change in investment cost ΔC. It should be noted that a reliability of 100% is never attainable.

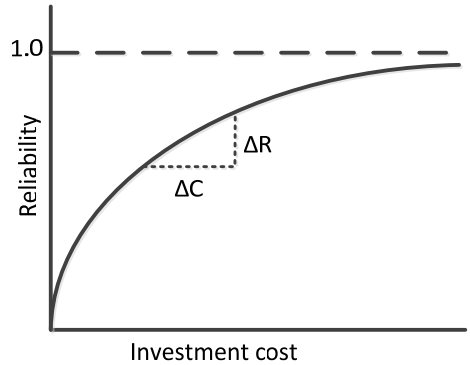

Figure 12. Incremental cost of reliability.

In practice, it is extremely difficult to find the true relationship between investment cost and reliability because of the complexity of the power system, the random nature of processes within the system (e.g., unscheduled component outages), and the subjectivity of outage costs. Reliability indices and measures are efforts to quantify reliability in the power system. In order to

deal with its complexity, power system reliability assessment divides the system into generation, transmission, and distribution. Probabilistic techniques are employed to plan for uncertainty in the load, component availability, and more recently, output power available from renewable energy sources.

The following sections are a summary of the indices and techniques that are most frequently applied in the reliability assessment of power systems.

A.1 Resource Adequacy

Reliability techniques and indices related to generation capacity are employed in power system planning where long time horizons (i.e., years) are considered. These methods help determine how much capacity is needed in order to meet expected future demand while keeping enough reserves to be able to perform corrective and preventive actions. The issue of whether installed capacity is sufficient to meet the electric load is known as *resource adequacy*. Descriptions of resource adequacy indices are presented next.

A.1.1 Loss of Load Probability (LOLP)

This index estimates the probability that the load will exceed the available generation during a given period. However, it gives no indication as to how severe the condition would be when the load exceeds available generation. For instance, two events can have the same probability of occurring (i.e., the same LOLP value), and the first one can belong to a generation deficiency of less than 1 MW, while the second one can belong to a generation deficiency of a few hundred MW. LOLP is expressed mathematically as:

$$LOLP = p(A - L < 0) \tag{1}$$

where A is the available capacity available to meet the system peak load L, and p denotes probability. Generally, LOLP is calculated by convolving the capacities and forced outage rates (FOR) of the installed generation fleet (Hsu, 1985). This produces a capacity outage probability table (COPT) that contains the probability of having outages of different MW levels. An example of a COPT is given in Table 3 for a system with 6 generating units and a forced outage rate (FOR) of 0.08 for each unit (NERC, 2011).

Table 3: Example of capacity outage probability for a 6-generator system with FOR of 0.08 for each generator

MW-out	MW-in	Probability	LOLP
0	300	0.60635500	1.0
50	250	0.31635913	0.00000026
100	200	0.06877372	0.07728587
150	150	0.00797377	0.00851214
200	100	0.00052003	0.00053838
250	50	0.00052003	0.00001835
300	0	0.00000026	0.00000026

Alternatively, a Monte Carlo simulation can be employed to calculate the LOLP of a system. Then LOLP can be expressed mathematically as:

$$LOLP = \frac{\sum_{i=1}^{N} S_e}{N} \quad (2)$$

where S_e is a simulation in which at least one significant event occurs. A significant event occurs when load and operating reserve obligations exceed resources or some event threshold limit. N is the number of years in the sampling period. Typically, there is one simulation for each hour of each year and LOLP is given as a percentage.

A.1.2 Loss of Load Expectation (LOLE)

This index is widely used when evaluating new generation scenarios in the planning process. It is generally defined as the average number of days on which the daily peak load is expected to exceed the available generating capacity (Allan, July 1992). Assuming a Monte-Carlo simulation is employed, LOLE in hours/year can be defined mathematically as (Billinton, 1991):

$$LOLE = \frac{\sum_{i=1}^{N} r_i}{N} \quad (3)$$

where i is the sampling year, r_i is the loss of load duration in hours and N is the number of years in the sampling period. LOLE and LOLP are directly related. Hence, LOLE has the same weakness as LOLP of providing no information about the severity of the condition.

A.1.3 Expected Unserved Energy (EUE) or Loss of Energy Expectation (LOEE)

This index is defined as the expected energy that will not be supplied due to those occasions when the load exceeds the available generation. Assuming a Monte-Carlo simulation is employed, EUE in MWh/year can be defined mathematically as:

$$EUE = \frac{\sum_{i=1}^{N} E_i}{N} \quad (4)$$

where E_i is the energy not supplied in MWh and N is the number of years in the sampling period.

A.1.4 Effective Load Carrying Capability (ELCC)

The ELCC is the contribution that a generator makes to overall resource adequacy. It quantifies the additional amount of load that can be served due to the addition of an individual generator (or group of generators) while maintaining the existing reliability level (Keane, 2011). ELCC is also known as *capacity value*. For conventional generators, the ELCC can be calculated based on their respective capacities and forced outage rates (FORs). These two are convolved using an iterative method to produce a capacity outage probability table (COPT) that indicates the probability of a given MW outage in the entire system.

Because wind capacity and FOR cannot appropriately describe the available wind power during peak load hours, the ELCC calculation must be modified to accommodate the uncertainty associated with wind. The IEEE preferred method to calculate the capacity value of wind consists of the following steps (Keane, 2011):

1. The COPT of the power system is used in conjunction with the hourly load time series to compute the hourly LOLPs without the presence of the wind plant. The annual LOLE is then

calculated. The LOLE should meet the predetermined reliability target for that period. If it does not match, the loads can be adjusted, if desired, so that the target reliability level is achieved.

2. The time series for the wind plant power output is treated as negative load and is combined with the load time series, resulting in a net load time series. In the same manner as step 1, the LOLE is calculated. It will now be lower (and therefore better) than the target LOLE in the first step.
3. The load data is then increased by a constant load ΔL across all hours using an iterative process, and the LOLE recalculated at each step until the target LOLE is reached. The increase in peak load (sum of ΔLs) that achieves the reliability target is the ELCC or capacity value of wind.

The use of Monte-Carlo simulation to evaluate multiple years is recommended in order to minimize the error due to inter-annual variation of wind.

Other methods for assessing the capacity value of wind exist such as using synthetic time series of wind in case there is a limited availability of historical wind data. One of the key factors in this case is to capture the correlation between wind output and load due to underlying weather conditions in the stochastic models of wind and solar plants.

A noniterative method to approximate the ELCC of wind that requires minimal modeling and is computationally inexpensive was proposed by D'Annunizo and Santoso (D'Annunzio, 2008). This method models a wind plant as a multistate unit that can exist in one or more partial capacity outage states C_j. Hence, a capacity outage individual probability table (COIPT) can be created with multiple discrete power levels (e.g., 0, C_j, $2 \cdot C_j$, $3 \cdot C_j$) up to the total capacity of the wind plant C_A. The probability p_j of a partial capacity outage state C_j is calculated by counting the occurrences when the power output is equal to C_A-C_j divided by the total number of power output data points. This can be expressed mathematically as:

$$p_j = \frac{\text{Number of ocurrences when power output is } C_A - C_j}{\text{Total number of power output data points}} \tag{5}$$

When a power output value falls between two discrete capacity outage states, it is counted as an occurrence for the highest value.

In addition to the wind plant COIPT, the method uses various load duration curves to determine the relationship between the LOLE of the system and an increase or decrease in the typical load demand. The new load duration curves are produced by taking the original system load curve and shifting it by a given number of percentages (e.g., –20%, 17.5%, –15%, …, 0%, ..., 15%, 17.5%, 20%). Mathematically, this can be expressed as:

$$L_c = L_t \pm c \cdot L_{t_{\text{pk}}} \tag{6}$$

where L_c is a new load duration curve, L_t is the typical load duration curve with peak load $L_{t\text{pk}}$.

Then, the LOLE is computed for each new duration curve. The resulting data points for LOLE as a function of peak load L_{pk} are fitted to an exponential function of the form:

$$LOLE(L_{\mathrm{pk}}) = B \times e^{m \times L_{\mathrm{pk}}} \tag{7}$$

Thus, an estimated value of m is found. The ELCC can be computed using:

$$\mathrm{ELCC} = \left[-\ln \left[\sum_{j=1}^{k} p_j \times e^{m \times (C_j - C_A)} \right] \right] \times \frac{100\%}{m \times C_A} \tag{8}$$

where C_j and p_j are the partial capacity outage states (MW) and corresponding individual probability, respectively. The nameplate capacity of the added unit is C_A.

Results from a case study shown in (D'Annunzio, 2008) showed that this method produced accurate results, within 3% of the ELCC value estimated employing the ELCC classical method described at the beginning of this section.

A.2 Transmission Reliability Indices and Measures

Reliability assessment of the bulk power system (i.e., transmission and generation) is divided into *resource adequacy* and *system security*. The previous section dealt with the issue of resource adequacy. This section now addresses the issue of system security, which refers to the question of whether the transmission system can move energy from generation to bulk supply points, while staying within operational limits and being capable of withstanding disturbances (Allan, Nov. 1992). In other words, the reliability of the transmission system must satisfy both dynamic conditions (i.e., withstanding a transient disturbance or small signal disturbance) as well as the static conditions (i.e., voltage, frequency, and thermal limits). Past performance indices applied to the transmission system include: system unavailability; unserved energy; number of incidents; number of hours of interruptions; number of voltage excursions beyond limits; and number of frequency excursions beyond limits.

As previously mentioned, the NERC has a very large number of standards that are employed by electric utilities in the mainland and that are oriented toward improving and assessing the reliability of interconnected electric systems.

A.2.1 Review of NERC Reliability Standards

The following is an overview of a subset of NERC's reliability standards. The standards mentioned below are organized alphabetically, as presented in NERC's complete set of *Reliability Standards for the Bulk Electric Systems of North America* (NERC, 2009).

Real Power Balancing Control Performance (BAL-001.1a)

This NERC reliability standard is aimed at keeping the steady-state frequency within defined limits. It defines the control performance standards (CPS1 and 2). In general terms, these control performance standards are statistical metrics of a balancing authority's ability to closely follow its demand in real time.

In interconnected systems, frequency deviations in combination with scheduled energy interchange values are employed to determine the mismatch between generation and load within balancing authorities.

***Automatic Generation Control (BAL-005-0.1b*)**
This standard establishes requirements for a balancing authority (BA) to calculate the Area Control Error (ACE) necessary to perform AGC. Examples of these requirements are maintaining regulating reserves that can be controlled by AGC, ensuring data acquisition for ACE calculation occurs at least every 6 seconds by having redundant independent frequency metering equipment, performing hourly error checks to determine the accuracy of control equipment, and periodically testing and recharging back-up power for control centers.

Operating Reserves (BAL-STC-002-0)
This standard provides a set of qualitative requirements that defines available operating reserves. These requirements are qualitative and do not set arbitrary operating reserve values, but give BAs a framework to determine reserve capacity necessary for reliable operation (e.g., operating reserves must be able to replace generation and energy lost due to forced outages of generation or transmission).

Cyber Security (CIP-002 to CIP-009)
These standards provide a framework on management and maintenance of cyber assets in power systems. These standards include functions such as identifying assets that are critical for managing the reliability of power systems and the vulnerabilities of those assets.

Telecommunications (COM-001-1.1)
This standard requires each BA to ensure proper functioning of their telecommunication facilities. Additionally, written operating procedures and instructions should be available to enable system operation when a loss of telecommunication capabilities occurs.

Emergency Operations Planning (EOP-001-0)
This standard requires each BA to “develop, maintain and implement a set of plans to mitigate operating emergencies”. An example of an emergency is a violation of the system operational limits. In such case, a balancing authority must have a plan to reduce load sufficiently to avoid system failures. Such a plan should include aspects such as communication protocols to be followed, controlling actions to resolve the emergency and staffing levels for the emergency.

Disturbance Reporting (EOP-004-1)
This standard requires BAs to record disturbances or unusual occurrences that result in system equipment damage, interruptions or jeopardize the operation of the system in order to study them and minimize the likelihood of similar events occurring in the future.

System Restoration Plans (EOP-005-1)
This standard requires BAs to develop plans and procedures to ensure that resources are available for restoring the electric system after a partial or total shut down. Restoration plans include items

such as training personnel, verification of restoration procedures through simulation and testing of black start units.

Plans for Loss of Control Center Functionality (EOP-008-0)
Each utility must develop a contingency plan to continue reliability operations in the event that its control center becomes inoperable. This contingency plan includes requirements such as procedures for monitoring and controlling generation, voltage frequency, and critical substation devices; and maintaining basic communication capabilities without relying on data or communication from the primary control center.

Documentation of Black start Generating Unit Test Results (EOP-009-0)
This standard addresses the testing of black start units and its corresponding documentation in order to ensure these units are capable of performing this function.

Transmission Vegetation Management Program (FAC-003-1)
This standard is aimed at minimizing outages and other events due to vegetation located on transmission right-of-ways and maintaining clearances between transmission lines and vegetation. It requires the transmission owner to have and update a formal transmission vegetation management plan.

Modeling and simulation of Interconnected Transmission System (MOD-010-0 and MOD-012-0)
The main purpose of these standards is to establish consistent models to be used in the analysis of the reliability of an electric system. It puts the burden of providing appropriate simulation models on power system component owners. For instance, it requires generator owners provide steady-state and dynamic modeling and simulation data to the regional reliability organization, which is the entity responsible for performing reliability assessments.

Aggregated Actual and Forecast Demand and Net Energy for Load (MOD-017-0.1)
This standard addresses the need for records of past and real-time load and demand-side management data. This data is necessary to forecast load and to perform future system reliability assessment. The standard gives some specifications on requirements such as “integrated hourly demands in MW for the prior year” and “monthly peak hour forecast demands in MW and Net Energy for load in GWh for the next two years”.

Reporting of Interruptible Demands and Direct Control Load Management (MOD-019-0.1)
This standard addresses the need for records and forecasts on interruptible loads and direct control loads to be employed in the system reliability assessment.

Verification of Generator Gross and Net Real and Reactive Power Capabilities (MOD-024-1 and MOD-025-1)
This standard makes generator owners responsible for ensuring that accurate information on real and reactive power capability of the units is available. This information is employed in the reliability assessment process. The standard includes requirements on the periodicity of data verification and reporting and the type of information to be reported.

System Personnel Training (PER-005-1)
This standard requires that each balancing authority, reliability coordinator and transmission operator use a systematic approach to training in order to address company specific reliability related tasks. Such training should be updated periodically in order to modify or add new tasks. The training program should also be evaluated periodically.

Analysis and Mitigation of Transmission and Generation Protection System Mis-operations (PRC-04-1)
This standard requires the transmission, distribution and generator owners to analyze protection system mis-operations and implement corrective actions to avoid similar events in the future.

Transmission and Generation Protection System Maintenance and Testing (PRC-005-1)
This standard requires transmission, distribution and generator owners to have a testing and maintenance program for protective devices in their systems. The program should include testing intervals and testing and maintenance procedures. Records of test results and maintenance should be maintained.

Other standards are parallel to this one, but applied to under-frequency load shedding (PRC-008-0) and under-voltage load shedding (PRC-011-0) equipment.

Under-frequency Load Shedding Performance Following an Under-frequency Event (PRC-009-0)
Transmission and distribution owners are required, under this standard, to perform an analysis of each under-frequency event to determine the performance of the under-frequency program. For instance, the cause of an under-frequency event, and load shedding set points and tripping times should be reviewed periodically.

Under-voltage Load Shedding Program (PRC-010-0-PRC-011-0, PRC-021-1-PRC-022-1)
Similarly to the under-frequency load shedding program, the under-voltage load shedding (UVLS) provides preservation measures to avoid voltage instability or collapse. The design and effectiveness of UVLS measures should be evaluated periodically (e.g., every 5 years). The UVLS equipment shall be maintained and tested periodically and data on the technical characteristics (e.g., breaking operating times, voltage set points and clearing times) shall be kept and updated.

Transmission Relay Loadability (PRC-023-1)
This standard requires that transmission operators adjust their relay settings so that they do not limit transmission capability of lines while still performing their protective actions appropriately. The standard gives quantitative guidelines for relay setting such as loadability at 0.85 per unit voltage and a power factor angle of 30 degrees.

Normal Operations Planning (TPO-002-0)
This standard addresses the need for planning in the power system at several time horizons, and communication of these plans or changes during operation between the different parts of the electric system. It requires the system operator to set plans for reliable operation through a

"reasonable" future time period; plans to meet unscheduled system changes using a single contingency (i.e., N-1) planning at a minimum; and to perform studies of next-day and current day conditions to determine system operating limits. Additionally, generator owners are required to communicate any changes in real output power capabilities and characteristics of their units.

Monitoring system conditions (TOP-006-1)

This standard requires that critical reliability parameters be monitored in real-time. These parameters include real and reactive power flows, line status, voltage, tap-changer settings, and system frequency. It also calls for weather forecasts and past load patterns to be used by the system operator in order to predict near-term load.

Response to Transmission Limit Violations (TOP-008-1)

This standard requires transmission operators to take immediate actions when system operating limits are violated. It also asks for the transmission operator to collect sufficient information and to use analysis tools to determine the cause of the violations in an effort to mitigate them.

System Performance under Normal to Extreme Conditions (TOP-001-0.1 to TOP-004-0)

These standards address the need for periodic simulation and assessment of system operation in order to ensure the reliability of the system in the long term. This assessment should be made annually using a near-term forecast (i.e., 1–5 years) and a long term forecast (i.e., 6–10 years). The purpose of these studies is to demonstrate that the system is able to perform up to a set of system standards for each of the following conditions:

- normal operation,
- loss of a single bulk electric system element,
- loss of two or more bulk electric system elements,
- and following extreme conditions.

The set of system standards can be found on page 950 of NERC, 2009 and is also attached in appendix A of this document. The transmission operator is required to upgrade or add components in order to meet future system needs and comply with the aforementioned standards.

Assessment Data from Regional Reliability Organizations (TPL-006-0)

This standard requires regional reliability organizations to provide system data, reports and system performance information necessary to periodically assess reliability and compliance. Examples of such data are resource adequacy plans; electric demand forecast and forecast methodologies; assumptions and uncertainties; supply-side resource information; and transmission system information.

Generator Operation for Maintaining Network Voltage Schedules (VAR-002-1)

This standard requires generator owners to provide reactive power control and voltage control necessary to maintain voltage, reactive power flows and resources within specified operating limits. Additionally, it mandates generators to notify system operators of any changes in reactive power capability.

A.3 Distribution Reliability Standards

Several metrics are employed when evaluating distribution system reliability. These metrics can be divided according to the length of the interruption and other data employed in their calculation as sustained interruption indices, load based indices, and other indices (momentary interruption). A momentary interruption refers to any interruption lasting less than 5 minutes and caused by the operation of an interrupting device such as circuit breakers. Consequently, a sustained interruption is any interruption lasting more than 5 minutes (IEEE, 2004).

Indices based on sustained and momentary interruptions take into account the number of customers affected by the interruption and the time it takes to recover from them. On the other hand, load based indices are those that focus on the load interrupted. The next section presents the most employed sustained interruption indices. Other distribution reliability indices can be found in Appendix B.

A.3.1 Sustained Interruption Indices

System Average Interruption Frequency Index (SAIFI)

This index indicates the frequency at which the average customer experiences a sustained interruption in the time interval under analysis (e.g., 1 year). Mathematically, this is given as:

$$\text{SAIFI} = \frac{\sum \text{Total number of customers interrupted}}{\text{Total number of customers served}} \tag{9}$$

System Average Interruption Duration Index (SAIDI)

This index indicates the average time an average customer experiences sustained interruptions in the time interval under analysis. Mathematically, this is given as:

$$\text{SAIDI} = \frac{\sum \text{Customer interruption durations}}{\text{Total number of customers served}} \tag{10}$$

Customer Average Interruption Duration Index (CAIDI)

This index indicates the average time that it takes to restore service after a sustained interruption in the time interval under analysis. In this index, customers with multiple interruptions are counted multiple times. It is expressed mathematically as:

$$\text{CAIDI} = \frac{\sum \text{Customer interruption duration}}{\text{Total number of customers interrupted}} = \frac{\text{SAIDI}}{\text{SAIFI}} \tag{11}$$

Average Service Availability Index (ASAI)

This index indicates the fraction of time that a customer has received power over a predefined period of time.

$$\text{ASAI} = \frac{\text{Customer hours service availability}}{\text{Customer hours service demand}} \tag{12}$$

The denominator is calculated by multiplying the number of customers served by the hours in the predefined period of time.

Customers experiencing multiple interruptions (CEMI$_n$): This is the ratio of total customers that experienced more than *n* sustained interruptions over the study period. It is expressed mathematically as:

$$\mathrm{CEMI}_n = \frac{\text{Total number of customers that experience} > \text{n sustained interruptions}}{\text{Total number of customers served}} \tag{13}$$

A.4 Economic Perspectives on Electric Grid Reliability

Most of the previous discussion relates to grid reliability measurement in the short-term with existing capacity and resources determined and unchangeable. Not only is the focus of the previous discussion short-term, it is completely supply-side centric. The issue of optimal reliability is not addressed. In contrast, much of the reliability economics literature pertains to the determination of the optimal level of reliability as desired by electric consumers. This literature is summarized in Appendix E. The incorporation of consumers' views of reliability value into the determination of optimal reliability requires that means are available to quantify this value. Three general techniques for quantifying consumers' values of reliability have been used:

- Market based methods attempt to examine actual customer behavior in response to various service options or investments in reliability to infer customer outage costs as evidenced by customers who sign up for non-firm service rates or install backup generation;
- After the fact measurement of actual outages that have occurred;
- Survey methods use customer responses to postulated outage scenarios to measure outage costs.

Many utility efforts to evaluate the outage costs of their customers have employed survey techniques most frequently and have attempted to elicit responses to:

- Direct costs: customer incurred costs of an outage of specified duration and advance warning time;
- Willingness to pay: customer outlay to avoid an outage of specified duration and advance warning time;
- Willingness to accept: payment received from utility to compensate for an outage of specified duration and advance warning;
- Revealed preference: question elicits customer response regarding specified combinations of increasing price (electric rates) and reliability (reduced outages);

Surveys are popular with utilities because they allow the utility to focus on the particular preferences of their customers and the unique outage and operating characteristics of their utility systems.[21]

[21] In contrast, survey methods are not popular with economists, although they are sometimes used. The concern is that consumers will respond to hypothetical questions by interpreting what they believe the interviewer wants to hear, or, based on their supposition of the reason the question is being posed. Or, consumers might try to game the outcome of the survey.

An integrated reliability planning conceptual model presented by (Burns and Cross, 1990) is shown in **Figure 13**. Total system cost, C_{total}, is the sum of system costs, C_s, and outage costs, C_o. Outage costs decline as the level of reliability increases. Correspondingly, system cost increases as investments to achieve increased reliability are made. Other things equal the optimal level of reliability occurs where the marginal increment to system costs needed to achieve an increment to reliability is equal to the incremental decrease in outage costs resulting from this level of reliability. While it is difficult to discern from the drawing this would occur at the minimum point of the total cost curve.

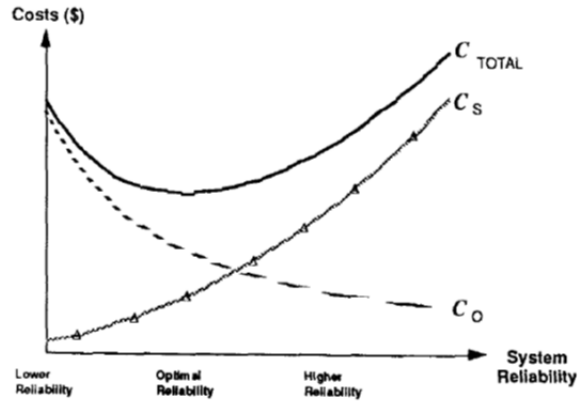

Figure 13. Hypothetical outage, system, and total cost as functions of reliability level.

Increased reliability requires increased capacity (larger reserve margin) which increases system costs; but outage costs decline with increased reserve margin, leading to the model shown. A simple mathematical model is derived from specification of these concepts and is manipulated to express marginal changes in expected unserved energy for marginal changes in capacity. This model is reproduced herein with some minor differences in notation.

An aspect of electric system reliability is that, while the regulated utility does not experience directly the costs of outages, its customers do. The utility regulatory commission internalizes the utility customers' costs; but it also internalizes the interests of the subject of its regulation—the utility. Hence, it is appropriate to develop a model that addresses both. An available model results in a reasonably simple framework within which to operationalize the determination of optimal reliability from an engineering-economic standpoint. This model employs the fact that reliability is functionally related to two dependent variables—system costs and outage costs. Reliability is a function of the amount of capacity on the system in relation to peak load. More capacity increases system costs but reduces outage costs. Because these dependent variables are inversely related an optimal capacity (from the societal point of view) can be determined by the capacity that equalizes incremental capacity costs and incremental outage cost reduction.

Total cost for electric service, TC, is determined by system costs, c_s , and outage costs, c_o. System costs include capacity, operation and maintenance, and all other costs to supply energy. Outage costs include costs customers incur during an interruption of service including lost output in all sectors as well as spoiled inventories. This yields,

$$\boldsymbol{TC} = \boldsymbol{c_s} + \boldsymbol{c_o} \quad . \tag{14}$$

Reserve is related to capacity of resources on the system and the peak system load at any point in time as follows

$$\mathbf{TC} = \mathrm{c_s} + \mathbf{c}\boldsymbol{R} = \boldsymbol{A} - \boldsymbol{L} \tag{15}$$

where R is the margin of reserves, A is the available capacity to meet load and L is peak system load where it is assumed that this is the annual peak load. R, A, and L are all considered random variables. A reliability event takes place whenever L exceeds A and R becomes negative. One measure of the frequency with which this happens (or could happen) is the loss of load probability (LOLP), a statistical measure, as implied by its name. This can be expressed as

$$\mathbf{TC} = \mathbf{c_s} + \mathbf{c}\boldsymbol{LOLP} = \boldsymbol{p}\,(\boldsymbol{R} < \mathbf{0}), \tag{16}$$

where p denotes probability.

When R becomes negative load shedding, brownouts, or blackouts occur all of which result in some quantity of unserved energy which we designate as u. Then,

$$\mathbf{TC} = \mathbf{c_s} + \mathbf{c}\boldsymbol{u} = \boldsymbol{E}\,(\boldsymbol{R} < \mathbf{0}), \tag{17}$$

where E is interpreted as expectation.

An operational rule of thumb has arisen through repeated use of the reserve margin as a static, point-estimate of system reliability. It is related to R as defined above but is not, strictly, a statistical measure. Using this point estimate we can redefine (17)

$$\boldsymbol{m} = \boldsymbol{a} - \boldsymbol{l} \tag{18}$$

where a is the total capacity of resources available to meet load and l is the system highest peak load. Equation (18) is closely related to (15) and is the point estimate drawn from a statistical distribution.

For the time period under consideration a is the total quantity of resources available and l is the maximum value of the load random variable. As in equation 4 above, u is related to and, more strongly, a function of m. Any addition of capacity or reduction of peak load increases the reserve margin. As m increases $u(m)$ decreases. Determination of the optimal$m$, m^* requires the inclusion of outage costs. One interesting feature of this model to note is that, in effect, two options can improve reliability: installing additional capacity and/or reducing peak load. These two options can be operated upon independently. Thus, optimality includes consideration of which of the two options is least expensive to implement. More than likely, at least up to some

level of reduction, load shifting from the demand side, has the prospect of being more cost effective purely because no additional, or very little, capital investment is required.[22]

We can now relate marginal capacity to the reserve margin, m. Designate s as marginal capacity cost per MW and let q denote a unit of outage cost represented in units of MWh. Then, assuming that the function u can be evaluated the following can be defined.[23]

$$\boldsymbol{d(m) = \partial u/\partial m, \forall\, m.} \tag{19}$$

In utility operational practice system operators will invoke emergency actions to avoid allowing operating reserves to fall to zero. These actions are initiated sequentially presumably in ascending order of cost and include shedding interruptible customers, voltage reductions, and customer appeals for load reduction. The final action is implementation of rotating outages. With each action, i corresponding expected unserved energy u_i can be interpreted as the energy "supplied" but the emergency actions, i. With I emergency actions, the ith being that of rotating blackouts, we can write

$$\mathbf{d(m) = \partial u/\partial m, \forall\, m}\boldsymbol{u} = \sum_{i=1}^{I} \boldsymbol{u_i} \tag{20}$$

Then the marginal reduction in outage cost defined in 6 becomes

$$\mathbf{d(m) = \partial u/\partial m, \forall\, mu} = \sum_{\mathbf{i=1}}^{\mathbf{I}} \mathbf{u_i} \tag{21}$$

where $d_i(m) = \partial u_i/\partial m$. This expression can be evaluated analytically or by a differencing approach.

The costs per unit of unserved energy, q_i, resulting from each emergency action i are required to complete the model. Then, at the optimal value of m designated as m^*, the following relationship holds,

$$\mathbf{d(m) = \partial u/\partial m, \forall\, m}\boldsymbol{s} = \textstyle\sum_{i=1}^{I} \boldsymbol{q_i d_i}\,(\boldsymbol{m}^{*)}. \tag{22}$$

As mentioned, the evaluation of $d_i(m^*)$ is available from a probabilistic reliability framework or from production cost models. One method of estimating electric consumer outage costs is described below.

With estimates of the value of service to customers, expected unserved energy (kWh) can be converted to dollar values. This process can be carried out explicitly in a production cost

[22] Selective load reduction can be implemented administratively with large electric consumers who could be induced to shift electricity consumption by a variety of financial payments. A fully integrated retail market would require the installation of significant additional infrastructure as is being discussed in Smart Grid programs.

[23] The function can be evaluated either analytically if it is assumed to be continuous and differentiable or can be assessed by "differencing" as suggested by the Burns and Cross.

modeling framework. Estimated values of service obtained by PG&E for their customer classes are shown in Table 4 Similar techniques could be employed by HECO to estimate customer outage costs by customer class.

Table 4: Direct Cost of a Summer Afternoon Outage with 1-Hour Notice

Customer Class	Average Outage Cost (1988 $/kWh)
Residential	4.05
Commercial	39.69
Industrial	6.78
Agricultural	3.53
System Weighted Average	18.63

Using the derived model and customer outage costs estimates as shown **Figure 13** the authors are able to derive a relationship such as that shown in **Figure 14**.

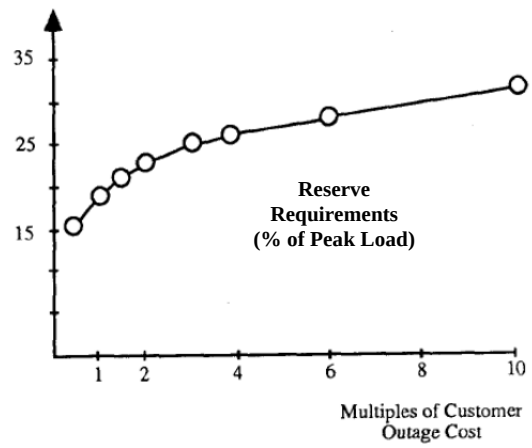

Figure 14. Variation of reserve requirements with respect to customer outage costs.[24]

Several interesting observations emerge from this paper. First, the significantly higher outage cost estimates for the commercial sector are notable as compared to all other customer categories that are somewhat closely grouped. While it could be expected to be somewhat higher, to have it be so much greater than the other sectors is surprising. Further, these data are from 1988 so clearly they would not reflect current economic values. Furthermore, there may be economic and social changes that affect the outage costs as seen by consumers. For example, our economy has become significantly more service than manufacturing based. And there might be significantly more home production (telecommuting) today than in earlier years that would increase these customers' views of outage costs. Finally, because we don't include the value of output of "true" home production services (child care, home schooling, grocery shopping, food preparation, cleaning services, etc.) in the national accounts, customers whose households consist of stay-at-home moms and dads may undervalue the cost of outages. These observations suggest that periodic administration of outage surveys and re-estimation of customer outage costs should be

[24] Burns and Cross, 1990.

performed by utilities. One of the points that Turvey made in his book on electricity economics is that service reliability must be the same for all customers. Burns and Cross also emphasize that their model is based on the assumption of equivalent reliability and cost for all customers. However, they point out that, because customers have distinct needs, a system of uniform power supply reliability is not the most economical means to meet individual needs. And we know from practice that all customers do not, in fact, receive the same level of reliability due to local system effects and differences.

A.5 Reliability Looking Forward

The new focus of reform efforts in the electric industry is to introduce and diffuse market competition as the means of allocating resources needed to supply electricity to customers. Much of this effort is currently focused on the supply side and the introduction and refinement of wholesale markets for electricity supply. Meanwhile, increased interest is evident toward development of retail markets and integration of these markets with the wholesale markets. Toward this end focus is increased on "demand response" as a form of load balancing. The understanding that both demand and supply have the potential to work in concert to balance system supply and demand and can help to improve capacity use for the supply system as well as keep electric rates (prices) lower for customers. The concurrent interest in Smart Grid that would provide the platform for customers to express their demand schedule for electricity supports this wholesale/retail market integration. A number of states including Texas and states in the PJM area have programs to elicit demand response. It is possible in this broadened environment to envision a market for reliability that would allow consumers to specify their requirements for power quality. These trends are congruent with observations and findings of market designers who have a goal of making electric power markets more efficient and effective.

(Hogan, 2005) has made the case that reliability levels must be worked out in a market in which every consumer has the option to participate. He brings the value of lost load (VOLL) back into the discussion. (Cramton and Stoft, 2011) have stated this more explicitly: "If reliability is not individualized then individuals know that they will not receive less reliability if they pay less for it, because they can be given less only if everyone is given less. Consequently, everyone will refuse to pay for collective reliability and all will attempt to enjoy a free ride." (p. 24.)

A.6 Integrated Reliability Index

NERC is developing an integrated reliability index aimed at increasing the transparency of the reliability assessment process. This integrated risk index (IRI) tries to include all aspects of the reliability assessment process and combine them to produce a single number ranging from 0 to 100. This single number indicates the historical risk found in a power system based on three main characteristics: major system events experienced, conditions that indicate if an adequate level of reliability has been attained, and compliance to reliability standards (NERC, 2012).

A review of the integrated risk index calculation proposed by NERC was performed by David Robinson (Robinson, 2011). His findings suggest that no connection exists between the metric proposed by NERC and changes made to a power system in order to improve reliability. This review is found in Appendix C.

APPENDIX B: GRID RESILIENCE METRICS AND DECISION-CONTROL ARCHITECTURES

As discussed previously, a key property of grid-resilience metrics is that their computation requires knowledge of both (1) the set of potential disruption events and (2) control actions taken to mitigate those potential disruption events. In other words, grid resilience is conditional upon both the disruption events and the actions taken to mitigate those events. As a result, a critical aspect of the proposed research involves concurrent development of the decision-control architecture(s) in which grid resilience metrics can be used. This linkage is independent of the specifics of whether the grid resilience metrics are being evaluated in operations or planning contexts, although the specific context is likely to dictate the fidelity of the decision-control architecture under consideration.

It is also in the context of decision-control architectures that grid-resilience metrics are truly differentiated from more traditional reliability metrics. Specifically, grid-resilience metrics allow for differentiation between and selection among possible likely future system states or trajectories (defined by the set of potential disruption events). This differentiation provides the foundation for advanced decision frameworks to support anticipatory mitigation, operate-through, and recovery situations induced by contingency and more extreme disruptions to grid infrastructure.

In the course of this discussion, we intentionally remain open to the specific algorithms and software packages providing specific functionality. However, we note that multiple instances of much of the necessary functionality presently exists within various research institutions across the U.S.

In addition to the set of potential threats defined (e.g., hurricane or other extreme weather events), the decision-control architecture must endogenously identify potential (short-term) contingency events such as line overloads. Such analysis proceeds in the context of an estimate of the uncertain system state, which in turn relies on the following core algorithmic technologies: state estimation, predictive forward simulation, and uncertainty quantification. The purpose of this analysis is to identify future situations (e.g., instabilities and subsequent component failures) that are likely to drive the system—if unmitigated—into low-resilience states.

Given a set of potential threats (both exogenous and endogenous), the next task is to develop a control algorithm that most efficiently moves the system to the state of maximum expected resilience. Such a system requires a description of the grid control elements (defining the action space) and the range of their feasible values in each projected trajectory. To address this challenge, one must leverage the mathematical formalisms of stochastic and robust optimization, and associated solution algorithms. Sandia and other national labs have successfully leveraged these decision-making paradigms in power-grid contexts of generation expansion and daily reliable operation to facilitate optimal decisions that maximize the expected value of a given performance metric across a scenario set.

Optimization can also play a role in scenario generation. Because the grid, due to growth of renewables and demand, is operating closer to feasibility boundaries, cascading collapses are more likely to cause major blackouts. Currently, the North American Electric Reliability Corporation (NERC) requires power systems to be "N-1 secure," capable of absorbing the loss of a single component. However, given that component failures are not necessarily independent events (e.g., in the case of hurricanes), it is often desirable to consider near simultaneous failures of multiple components (N-k events). Identification of high-impact failure combinations is an extremely challenging combinatorial optimization problem. Recent state-of-the-art approaches avoid explicit enumeration and screening of contingency states by performing worst-case interdiction analysis via solutions of bi-level programs. The scalability of these approaches to practical-sized systems, however, remains a challenge. There exist preliminary successes in this area, scaling to systems with several hundred buses. However, further research is required to achieve scalability to full-sized industrial systems.

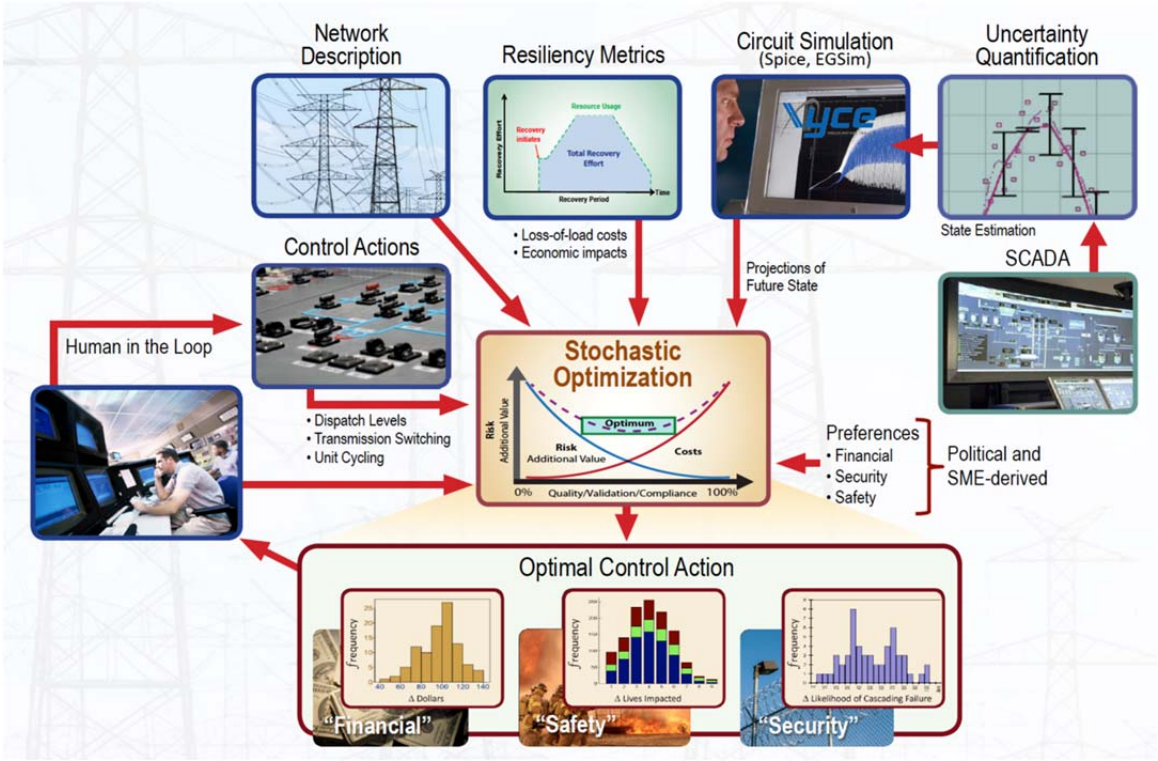

Figure 15. Decision-control architecture for grid resilience.

In **Figure 15**, we summarize the key components of decision-control architectures for grid resilience, emphasizing the interrelationships between the various components. The functionality of these components is now summarized, and their connections are described at a high level.

- **Network (System) Description.** The network or system description provides a full specification, at the necessarily level of fidelity, of all system components. Examples include transmission-line electrical properties, generator performance characteristics, and relay states. While typically viewed as static, the control architecture (specifically the stochastic optimiza-

tion component) can initiate actions that modify the network/system state (e.g., switching a line off).

- **Resilience Metrics.** Resilience metrics are mathematical computations that quantify some aspect of system resilience, taking as input (1) parameters relating to the risk and consequence of particular system failures and (2) a specification of the current system state. In practice, a resilience metric is typically a real-valued function. However, we leave open the possibility that the function is not closed-form, requiring a simulation or other algorithmic process to compute.
- **Uncertainty Quantification.** Due to limited precision and availability of sensor measurements, system state is necessarily uncertain. When operating at performance margins, accounting for this uncertainty is critical—ignoring the fact can lead to component failures, which in turn can lead to subsequent failures and system degradation. The purpose of the uncertainty-quantification component of a resilience decision-control architecture is to (1) perform state estimation under uncertainty (e.g., using supervisory control and data acquisition [SCADA] systems), and (2) to couple the state estimate with forward simulation (e.g., PSLF or SPICE), in order to predict subsequent system states—with corresponding probabilities. These predictions comprise a set of scenarios, above and beyond any predefined sets of disruption events. The latter are not a function of state estimation, but are rather exogenously specified. In contrast, the former are endogenous, and rely on current system state estimates.
- **Control Actions.** This component represents a description of the set of control actions available to a system operator, whether human or algorithmic. Examples of control actions include generator dispatch level adjustment, line switching, renewables curtailment, and allocation/ positioning of system-restoration resources.
- **Stochastic Optimization.** The above components either generate or directly serve as inputs to the stochastic optimization component, which rigorously determines those near-term control actions that will maximize system resilience. This maximization proceeds in the context of either one resilience metric or a weighted combination of multiple resilience metrics. The maximization is conditioned on the set of input scenarios, both those resulting from uncertainty quantification and predefined disruption events. Stochastic optimization will identify a single set of (nonanticipative) control actions that maximizes resilience across the set of scenarios, either in terms of a simple expectation or a risk-oriented aggregate measure. Thus, resilience should be viewed as a histogram, due to the consideration of a diverse set of scenarios, each representing a possible future system state.

It is important to observe that while we believe the identified components should be present in any such architecture, the fidelity is allowed to vary dramatically, depending on the usage context. For example, coarse-grained models of stability and other dynamics can be leveraged in long-term planning models, which can similarly rely on a small set of predefined disruption events. Similarly, many of the components can be simply simulated (e.g., the products of state estimation). Finally, we note that the resilience quantities at the bottom of the figure reinforce key properties of the metrics: their distributional nature (due to dependence on disruption

scenario and uncertainty in system state) and the presence of disparate and likely competing metrics.

APPENDIX C: ELECTRIC USE CASE

This appendix details an analytic example of resilience analysis for the electricity grid, using the proposed resilience framework and associated metrics. The presentation is organized around a series of analysis use cases, ranging from a baseline resilience computation for an existing system, to comparison of alternative investment portfolios to enhance resilience, to optimization of investments for enhanced resilience given a fixed budget.

C.1 Illustrative Test System

To illustrate the use of our resilience framework and metrics in the context of the electricity section, we consider a simple, well-understood, and widely used model of an electricity grid: the IEEE 118 bus test case. This model—available from http://motor.ece.iit.edu/data/itscuc—consists of 91 loads, 54 generators, and 186 lines. A high-level schematic of this test system is as follows:

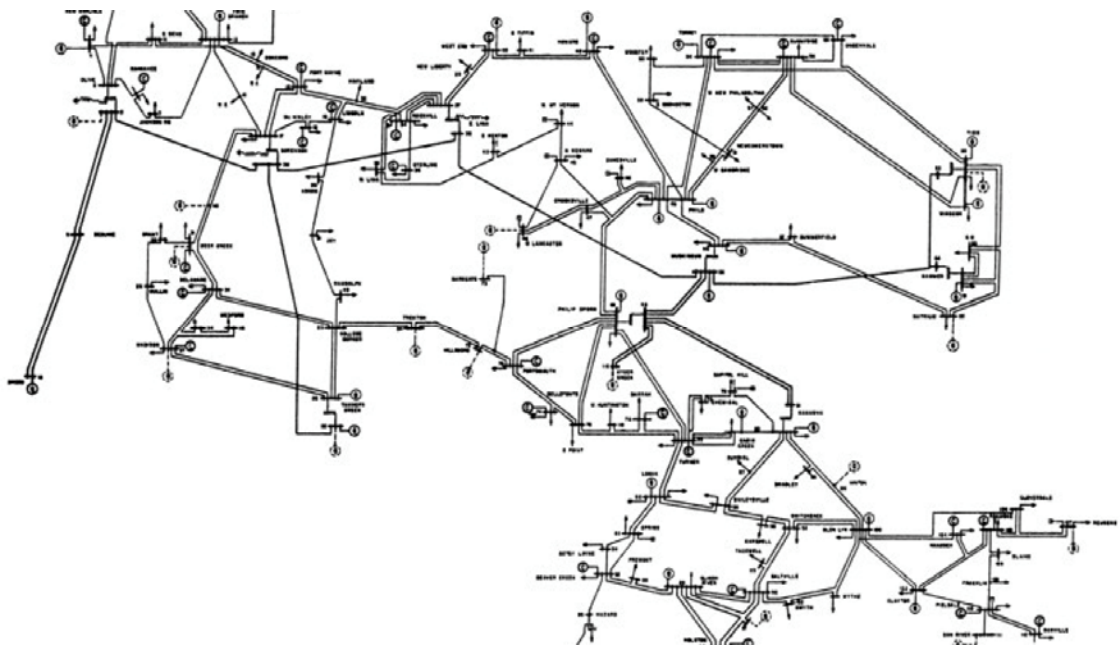

Figure 16. IEEE 118 Bus Electrical Test System.

The resilience framework assumes the availability of a systems operations model. For the 118 bus test case, we consider a standard security-constrained unit commitment model with economic dispatch, representing system reliability operations for a period of 24 hours. Network physics are approximated using DC optimal power flow models.

In terms of mathematical constructs, the unit commitment operations model is expressed as a mixed-integer linear optimization model, with algebraic constraints and objectives. This decision formalism contrasts with simulation formalisms, in that it allows for efficient global optimization of operational models. Specifically, the use of such algebraic optimization models facilitates automatic determination of optimal investment portfolios, subject to fixed, pre-specified budgets.

C.2 Defining Threat Scenarios

A key concept in the proposed resilience framework is the following: an infrastructure is designed to be resilient to a specific set of possible disruptions. In other words, the set of possible events is finite and pre-specified, at an appropriate level of abstraction. A number of methodologies can specify such events. However, for purposes of simplicity, we adopt the notion of a scenario tree. A scenario tree is a decision tree that specifies, via branching, the nature of the range of disruptions to which we are designing an infrastructure to be resilient to.

For a posited system, we consider three classes of high-level threats: a hurricane, an earthquake, and a terrorist incident. These three event classes correspond to the first branches from the root on the tree, yielding three "children" nodes. In general, probabilities can be assigned to each specific event class—assuming sufficient information is available for their estimation. More commonly, the event classes reflect all-hazard events, such that the probabilities are treated as uniformly distributed. In a real system, the high-level threat scenario identification process is expected to be an output from an iterative and highly interactive stakeholder-driven process. In this electricity use case, the first stage of the scenario tree can be depicted as follows:

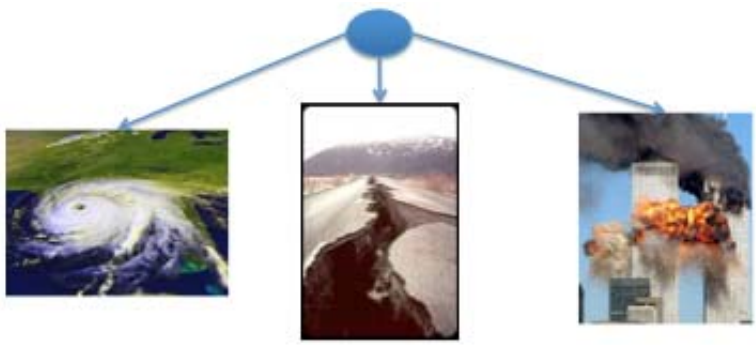

Figure 17. Threat scenario tree representing three "children" nodes.

Given a high-level threat specification, the next stage in the scenario analysis process is to further refine the description of each specific threat. For illustrative purposes, we focus on the hurricane scenario. In the case of a hurricane and other natural disruption events, significant historical information and forecast models can be used to guide specification of possible realizations of the general threat. For example, the scenario tree node representing a hurricane is expanded as follows:

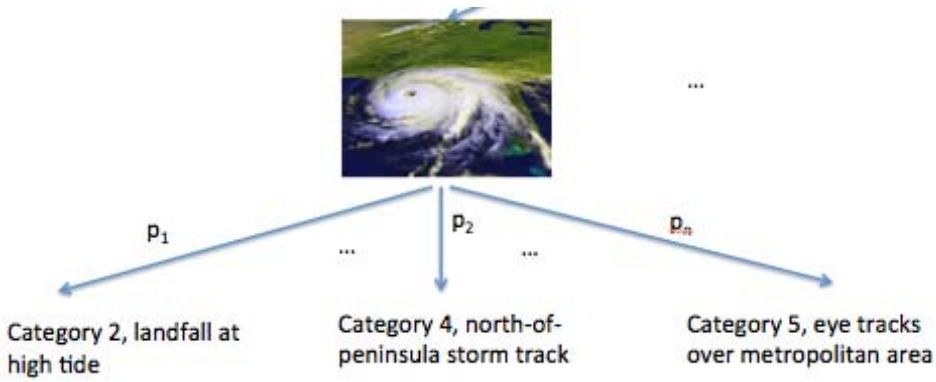

Figure 18. Detailed expansion of the hurricane threat.

Note that at this point in scenario analysis, probabilities for specific event realizations are likely to be available, or at minimum, relative weightings of likelihood.

Finally, each realization of an event must be translated into physical damage of the infrastructure system under consideration, e.g., the electricity grid. Pictorially, we illustrate this process for one of the hurricane events as follows:

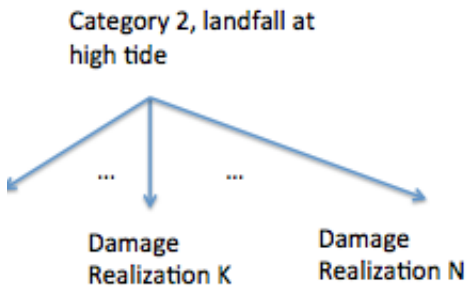

Figure 19. Damage realizations for a Category 2 hurricane, landfall at high tide.

For a real-world system, system experts should be consulted to estimate the damage to equipment given the occurrence of the threat. In the context of the IEEE 118 bus test case, we arbitrarily define system damage under this specific hurricane event realization as follows. For generation, we sample the number of distinct failures from a normal distribution, with mean 20 and standard deviation 5; the failures are then allocated uniformly and randomly to the generation fleet. We follow an analogous process to simulate damage to lines, using a normal distribution with mean 40 and standard deviation 7. These damage profiles are intended to be strictly notional. As with high-level threat scenario identification, actual damage realization profiles will need significant domain expertise and stakeholder involvement in order to be accurately specified.

C.3 Specifying Consequences of Loss of Delivery

In the context of the electricity grid, loss of power—typically quantified as MW hours of load shed—is a very indirect method to quantify the true consequence associated with loss of delivery. For purposes of resilience analysis, more salient metrics quantify aspects of safety, security, or economic impacts. In this analysis of the 118 bus test case, we consider economic losses calculated using hypothetical relationships, e.g., due to industrial facilities being disrupted. In

order to translate from loss of load at each bus to consequence, we introduce piecewise linear transformations—an example of which is given as follows:

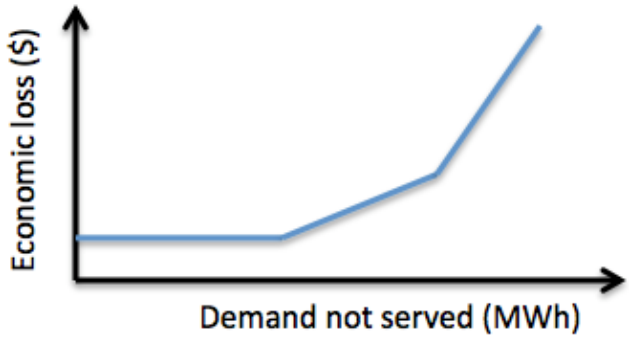

Figure 20. Example function relating MW not served at a bus to economic loss.

The particular transformation was chosen to reflect the typically nonlinear nature between loss of delivery and consequence. In reality, specification of this transformation requires deep knowledge of both the system under consideration *and* any associated dependent infrastructure.

C.4 Baseline Resilience Assessment

The first resilience analysis of the grid system involves computation of a baseline resilience value. The intent of this analysis is to demonstrate the initial use of a resilience metric for any infrastructure, which is to establish a rigorous and quantifiable description of system resilience. Without such a baseline measure, it is difficult to assess the benefit conferred by any proposed investments to improve system resilience.

We consider a hypothetical hurricane event, and sample 100 realizations of potential damage using the distributions of damage to generation and line resources introduced above. For each scenario, we compute a minimal-cost commitment and dispatch, which additionally minimizes loss of load. Given this dispatch, we then compute the cumulative economic losses incurred due to loss of service, using the piecewise linear transformations described previously. We assume no recovery is possible in the short term (i.e., 24 hours), such that generators and lines that have failed remain disabled for the scheduling horizon.

A histogram of cumulative economic losses incurred across the sampled scenarios is shown below:

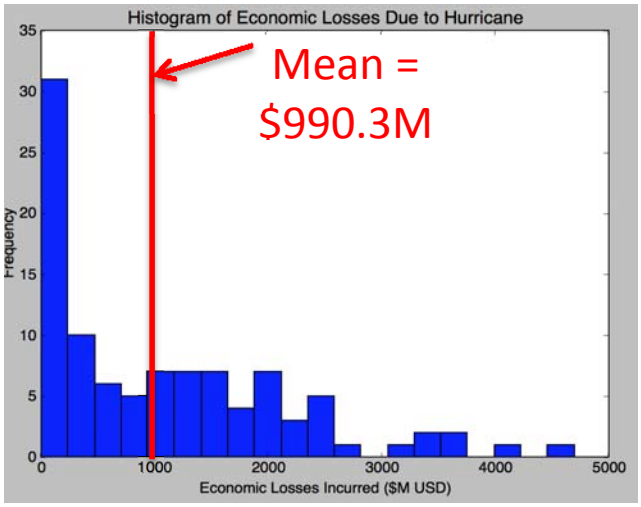

Figure 21. Resilience metric for the baseline system.

While many scenarios yield minimal economic losses, there are a nontrivial number of scenarios in which the economic loss is significantly larger than the mean of $990.3M. Translation from the distribution to a single summary statistic can proceed in a variety of ways, including a simple mean (as shown in the figure) or tail-oriented statistics such as Conditional Value-at-Risk (CVaR).

Beyond establishing a baseline resilience quantity, it is possible to simply operate the system to directly minimize consequence—as opposed to an economic dispatch. Under this paradigm, we are able to largely mitigate the expected consequences and VaR associated with the hurricane event. We graphically show the impact of shifting operations from an economic dispatch (minimizing operating cost) to a resilience dispatch (minimizing expected resilience) as follows:

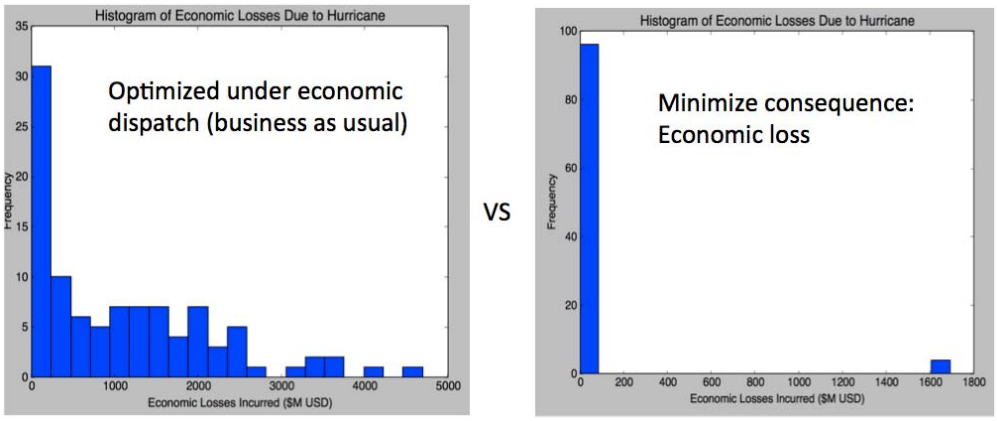

Figure 22. Resilience metrics for the same system under economic dispatch (left) and resilience dispatch (right).

In other words, by exposing consequence as a resilience metric and directly optimizing against this metric, it is possible to significantly reduce the consequences associated with a posited event.

C.5 Restoration and Recovery Analysis

Another aspect of resilience quantification relates to the time and costs associated with system recovery and restoration. To illustrate resilience analysis using these concepts, we augment the baseline 118 bus test case with a recovery and restoration model. This process is modeled as occurring over a three day period following the initial event. We assume a fixed budget for recovery and restoration resources, and impose the following:

- Five crews are available, 3 for line restoration and 2 for generator restoration
- Each crew requires 3 hours to repair a line
- Each crew requires 18 hours to repair a generator
- Lines are repaired in a random order
- Generators are repaired largest-to-smallest (in terms of capacity)

Mirroring the previous baseline resilience analysis methodology, we compute restoration and recovery costs for the associated disruption scenarios. The resulting histogram of cost is as follows:

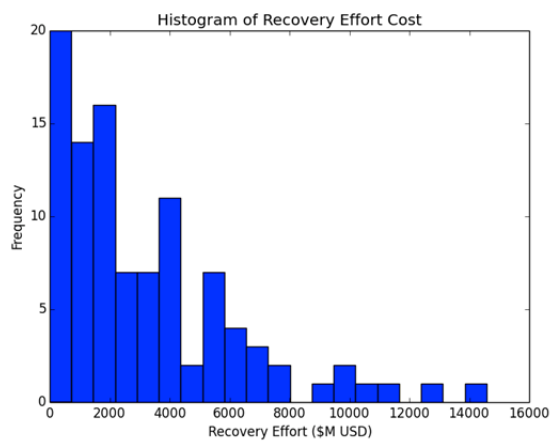

Figure 23. Resilience metric focused on restoration and recovery.

As with the analysis of economic losses incurred, methods for reducing the recovery and restoration distribution to a single metric include both expected value computations and tail-oriented statistics.

C.6 Investment Analysis

Given a baseline resilience analysis for a particular infrastructure system, the next logical step is to assess how different investment portfolios are likely to improve system resilience, and by how much. To illustrate the execution of this type of analysis, we consider the assessment of two competing investment options for the modified 118 bus test case. In Option A, engineers propose to build flood walls around generators with greater than 180MW capacity; this represents approximately 20% of the thermal fleet. This posited strategy—a proxy for protection against flooding—costs $100M total, $9.1M apiece for each of the 11 affected generators in the system. In Option B, engineers propose to bury high-capacity lines, specifically those with thermal limits

exceeding 250MW; this represents approximately 5% of the lines in the system. This posted strategy—a proxy for protection against high winds and tree faults—also costs $100M, $4M apiece for each of the 25 affected lines.

Re-running the baseline analysis with the additional system protections in place, we obtain the following:

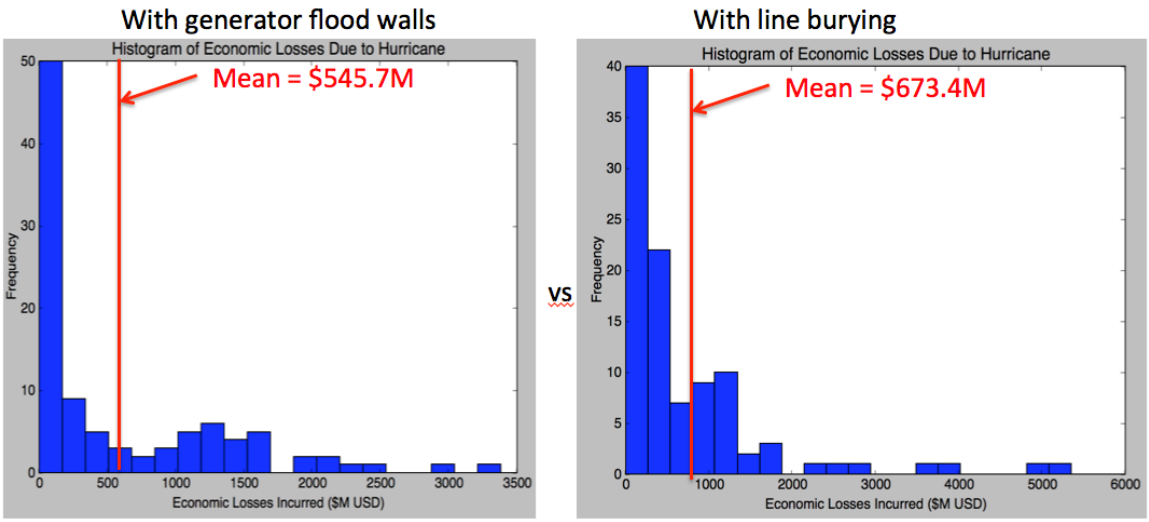

Figure 24. Resilience metrics showing improvements over baseline by adding flood walls (left) and burying electrical cables (right).

Note that both investment options reduce consequences, relative to the baseline mean of $990.3M. However, Option A yields a more significant reduction, and further admits fewer high-consequence events. Overall, the intent of this example is to illustrate the use of the proposed resilience framework and metrics to rigorously assess the relative benefits of proposed investment options—a critical step in (for example) rate case justification.

C.7 Advanced Planning

An alternative to evaluating competing investment portfolios is to simply determine the optimal investment portfolio directly, i.e., the portfolio that maximizes the increase in resilience (decrease in consequence) subject to a fixed budget constraint. In the case of this electricity use case example, this capability is enabled by the availability of the operations model as an algebraic mathematical optimization model. To illustrate this type of analysis, we expand the investment analysis scenario as follows. First, we assume a total budget of $100M, and respective hardening costs as previously specified—$9.1M per generator, and $4M per line. However, we introduce decision variables into the operational model that allow the optimization to determine which assets are hardened, and in what mix, while requiring that expenditures do not exceed $100M.

The resulting histogram of economic impacts incurred is as follows:

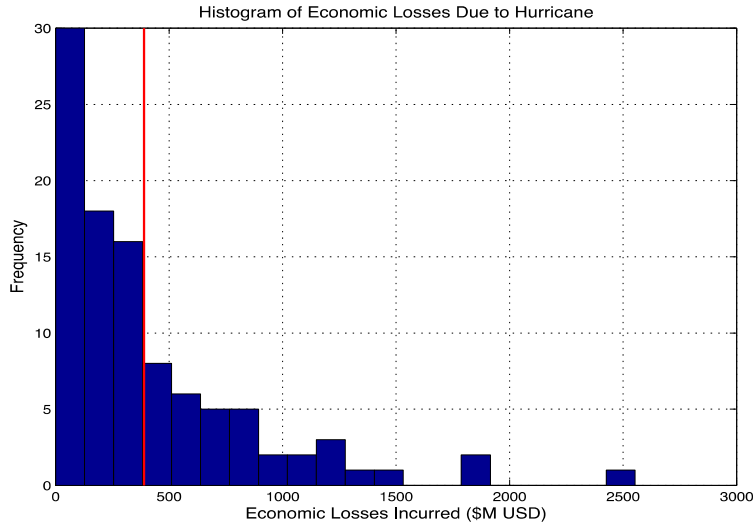

Figure 25. Resilience metrics showing improvements over baseline by combining the addition of flood walls with buried electrical cables. The optimization algorithm selected the mix of each, constrained to $100M total.

In this analysis, we minimized the mean (expected) economic loss incurred across 100 sampled scenarios of realized damage for the posited hurricane event. The graphic indicates that the resulting investment portfolio (which includes a mix of generator flood walls and line burying) outperforms both Option A and Option B, in terms of both reducing the mean impact and admitting fewer very high-consequence events, although the investment was equivalent.

C.8 Summary

In this appendix, we have illustrated the resilience framework and associated metrics in the context of the electricity grid. Using a standard test system, we illustrated how threat scenarios and consequences can be specified. Given this context, we demonstrated a range of resilience analysis for the test system, ranging from a simple yet critical baseline resilience computation to automated deterministic of investment portfolios to maximize improvement in system resilience.

APPENDIX D: PETROLEUM USE CASE

This use case demonstrates one way to use the resilience framework to identify potential options to increase resilience and measure the increase in resilience due to implementing these options. Specifically, we calculate the increase in resilience gained by re-engineering two major transmission pipelines to decrease down time after a large scenario earthquake in the New Madrid Seismic Zone.

D.1 Scenario Description

The New Madrid seismic zone (NMSZ), stretching along the Mississippi River Valley from southern Illinois to Memphis is the site of some of the largest historical earthquakes to strike the continental United States.[25] The last of these very powerful earthquakes occurred in the winter of 1811–12 when four major shocks occurred over a period of 48 days. This area was only sparsely populated at that time, so damage to buildings and other structures was limited. A repeat of that earthquake event today would not only cause a human catastrophe in the region directly damaged, but would also cause extensive damage to our nation's critical infrastructures. The US Geological Survey estimates a 7%–10% chance of an 1811–12 magnitude earthquake occurring in any 50-year period.[26]

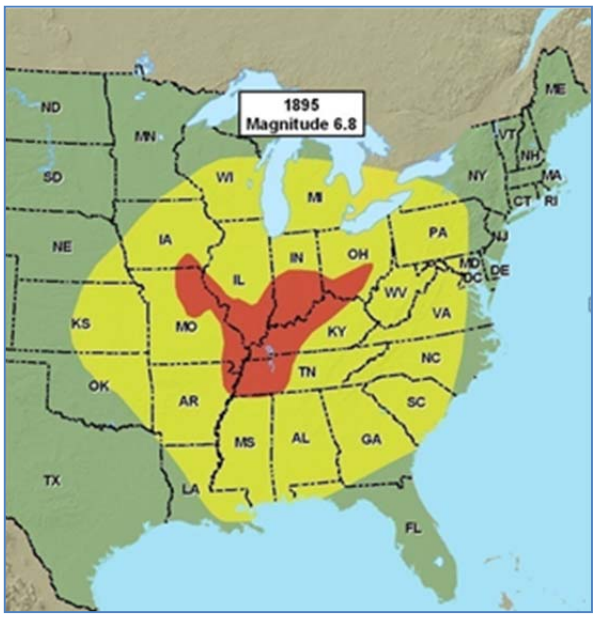

Figure 26. The region impacted by an 1895 earthquake. The red region indicates the extent of severe damage to structures. The yellow region indicates the area over which shaking due this earthquake was felt by observers.

[25] J. Gomberg and E. Schweig, "Earthquake Hazard in the Heart of the Homeland," US Geological Survey Fact Sheet FS06-3125, 2007, (http://pubs.usgs.gov/fs/2006/3125/pdf/FS06-3125_508.pdf). See footnote 1.

[26] Ibid.

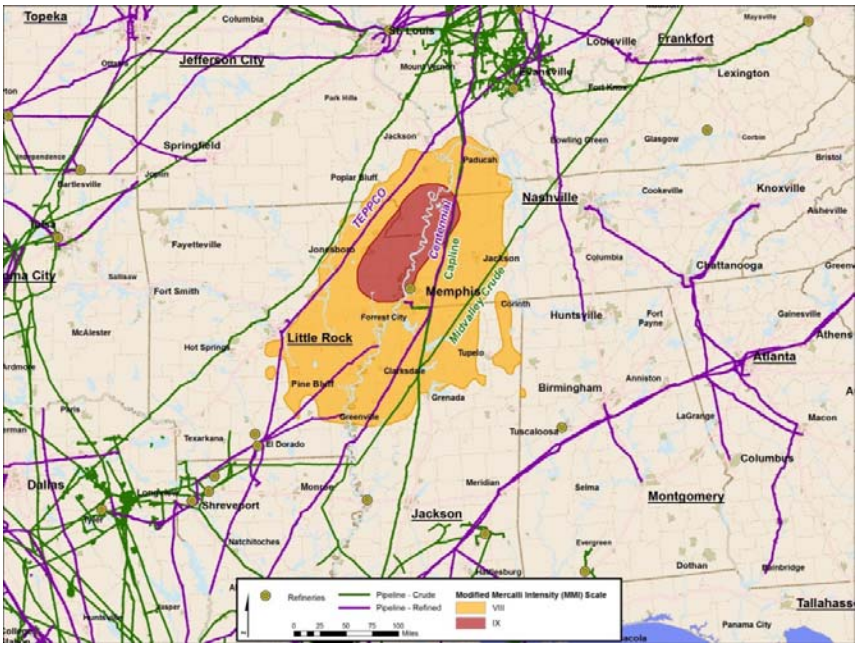

Figure 27. There are four major transmission pipelines that could be severely damaged by a New Madrid earthquake. Two of these, the Capline and the Midvalley carry crude oil. The TEPPCO and Centennial carry refined products. The shaded regions represent the shaking intensity of the scenario earthquake used for this analysis. Yellow and red indicate shaking intensities of VIII and IX, respectively, on the Modified Mercalli Intensity Scale.

D.2 Evaluating Resilience

Three models, a repair model, a petroleum network model and a consequence model were used to evaluate the resilience of North American petroleum infrastructure to a New Madrid earthquake. Results of the repair model serve as input to the network model, and results of the network model in turn feed the consequence model. Prototypes of two of the models were developed for this demonstration of the framework. An existing network model of North American petroleum infrastructure, the Sandia National Transportation Fuels Model (NTFM) was used to simulate the system level response to the scenario disruption. Notably, other models could be substituted for any or all of the models we used to demonstrate the framework for the petroleum case.

D.2.1 Repair Model

This prototype model estimates the cost and time of repair of a single infrastructure component given assumptions about the extent of damage, the level of preparation, and the availability of resources. To accomplish these estimates, the model represents role of logistics factors that could alter schedule or costs. For example:

- Delays in receiving materials and equipment
- Labor constraints
- Establishing field offices and communications
- Housing and transportation
- Level of material stockpiles and pre-planning

The Repair Model was developed using System Dynamics (SD) methodology. SD models are commonly used for analysis of supply chains and managing large projects. Models of this type are useful for these applications because they include the dependencies of one part of a project on another, constraints on the supply of materials or labor, time delays and the accumulation of costs.

The Repair Model simulates four phases of the repair process: Assess, Obtain Materials, Repair, and Test & Certify. The desired target schedule was established by providing estimates on each phase's duration and resource needs. Given sufficient resources, the model would indicate that the repair would be completed on schedule and budget. If however, there were unexpected delays or shortages of resources, the model would indicate a longer repair time and larger cost.

One use of this model is to evaluate how investments in preparation or infrastructure hardening could result in decreased repair times and costs. By providing estimates of repair times to the network model, each proposed investment can be associated with a value of the consequence metric.

For this demonstration, the Repair Model used nominal values and aggregated resource units. A resource unit represents labor, materials, and equipment needed to accomplish each phase. The resource usage set relatively among the four phases. In an actual use, experienced repair professionals would contribute estimates of the time and resources needed to complete the repair.

Figure 28 and **Figure 29** show examples of the two main outputs of the Repair Model, the calculated cost of repair versus time, and calculated functioning capacity of damaged component versus time. The repair time is the interval from the time of the damage until the time the functioning capacity returns to normal (indicated as 1 on the graph in **Figure 29**).

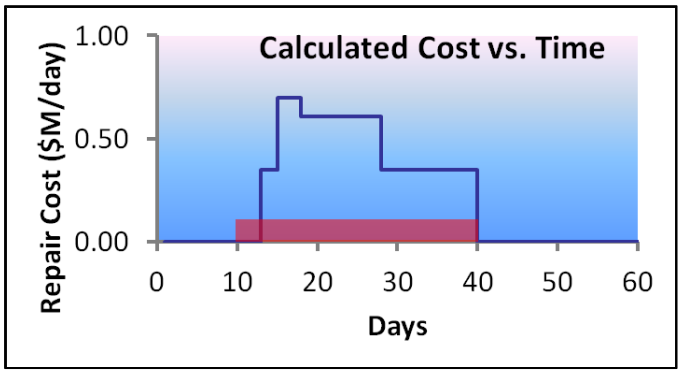

Figure 28. Repair Cost trend calculated with the Repair Model.

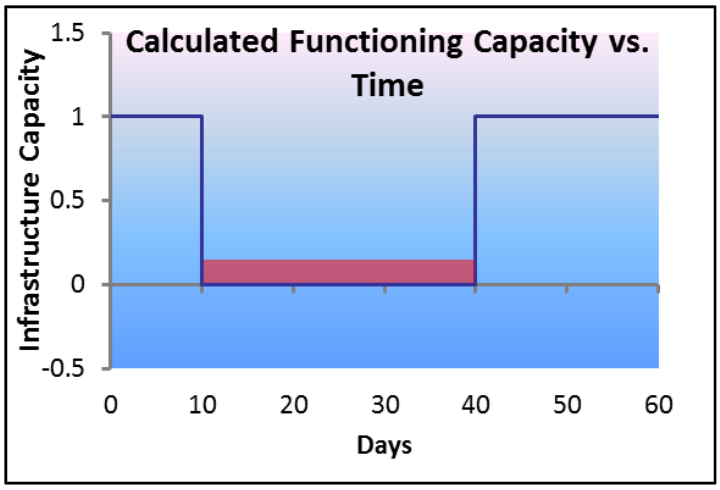

Figure 29. Functional Infrastructure Capacity of damaged component calculated with the Repair Model.

D.2.2 Network Model

The petroleum network model estimates the availability of transportation fuel in the event any component of the national fuel supply chain is damaged or disrupted. The portion of the fuel supply infrastructure represented by the model spans from oil fields to fuel distribution terminals. Different components of this system (e.g., crude oil import terminals, refineries, transmission pipelines, and tank farms) can be disrupted, and these disruptions can cascade through the system.

In order to evaluate system-level resilience, it is necessary that the model simulate all of the major system attributes or behaviors that increase resilience. Market-driven resilience attributes represented by the model algorithms include:

- Re-routing shipments
- Drawdown of inventory
- Use of surge capacity
- Increasing imports
- Reducing consumption

It is also necessary that the simulation results be constrained by connectivity of the system and capacity of individual system components:

- Pipeline flow
- Refinery throughput
- Tank Farm storage
- Import terminal throughput

This model represents the transportation fuel system as a network consisting of tank farms, refineries and terminals (the nodes of the network), and the pipelines that connect the nodes (the links of the network) (**Figure 30**). Sources of crude oil to the network are nodes that represent

either collections of oil fields (called geologic basins) or water terminals at which imports of crude oil are received. The close correspondence of the elements of the actual and model networks allows analysts to simulate damages to the network, and the resulting fuel availability impacts, at a reasonably high level of spatial resolution.

The model incorporates algorithms that seek to minimize fuel shortages while balancing mass and not exceeding capacities. There are some important model assumptions:

- Includes transmission system (pipelines, water*), but not distribution (trucks)
  - For example, the model does not know that fuel can't be delivered because roads are damaged
- Market behavior is based on fuel availability
- No hoarding behavior (by consumers or suppliers)
- No price increases until inventories decline
- Desired consumption of fuel is not decreased by damage to other infrastructures

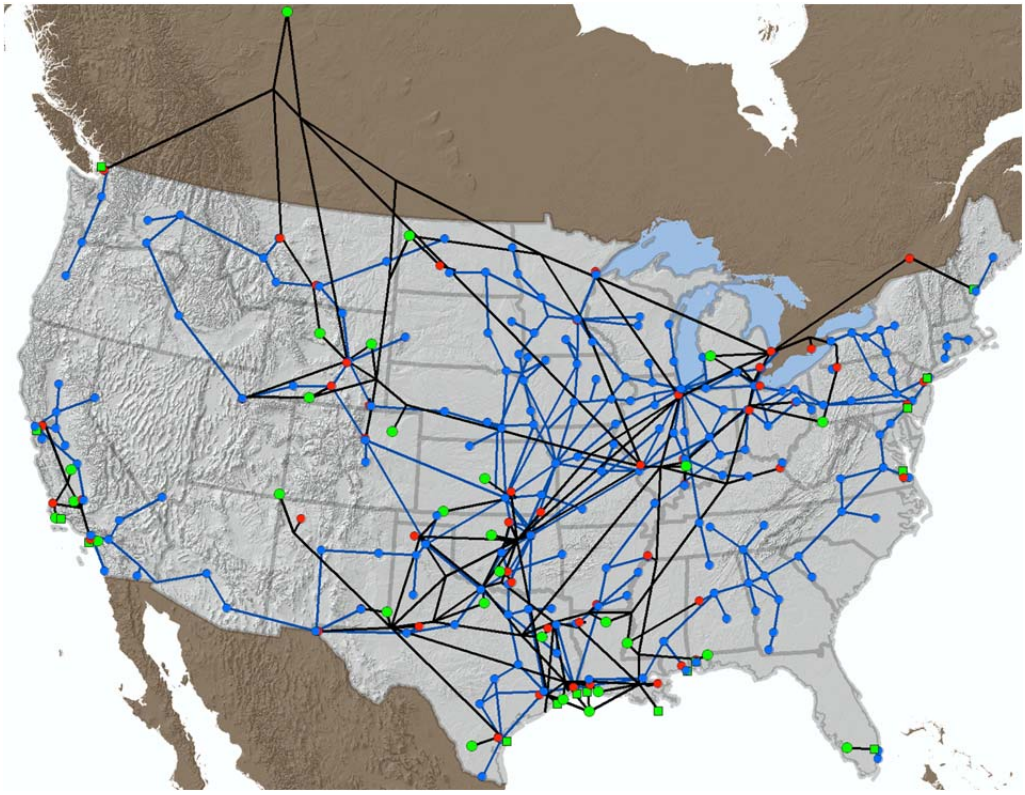

Figure 30. The transportation fuel network model.

An output of the model is the time history of fuel consumed from each distribution terminal. **Figure 31** shows for example, simulated fuel consumption at selected terminals that were impacted by a simulation of a 30-day disruption of the four transmission pipelines. The aggregate decrease in fuel consumed due to the scenario disruption is a measure of how well the system

performs its primary function while under stress. It is therefore the performance indicator calculated by the network model.

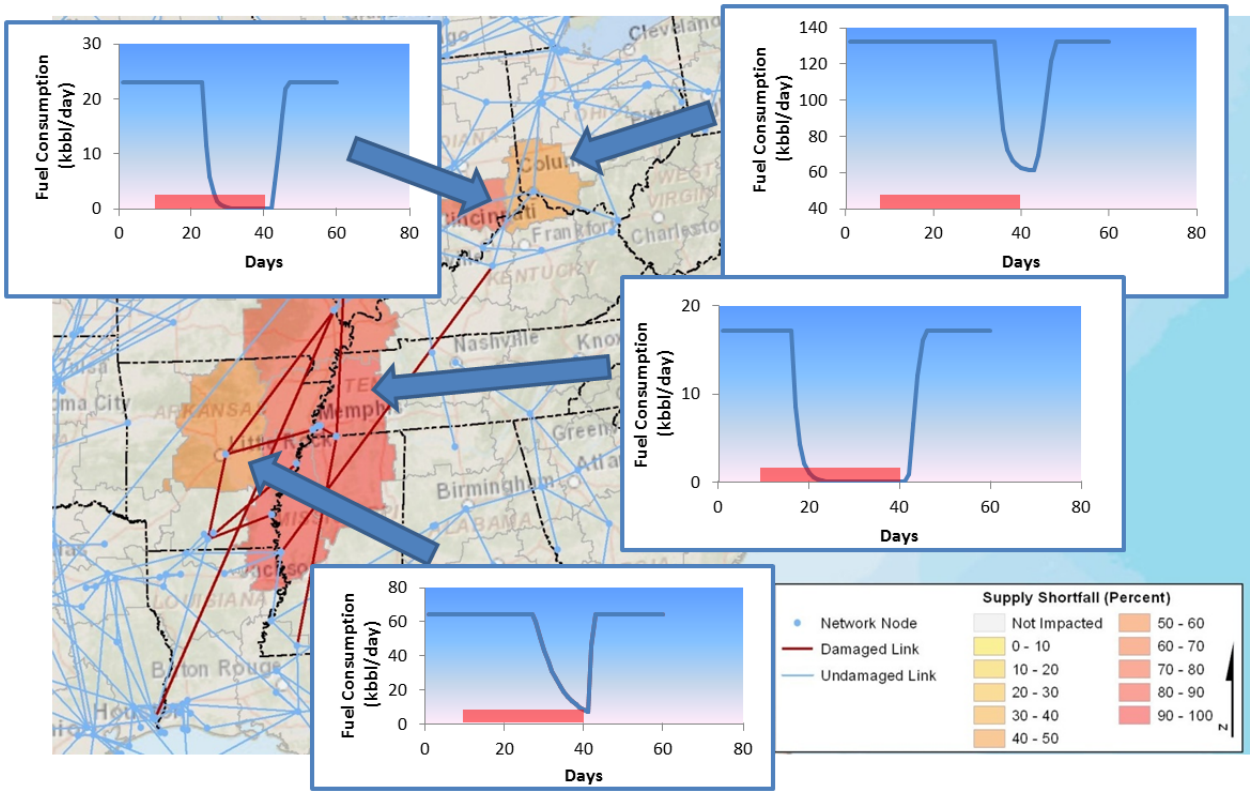

Figure 31. Simulated fuel consumption at terminals during 30-day disruption.

D.2.3 Consequence Model

In this example, we use the added fuel cost to consumers as a consequence metric due to the scenario disruption. There are three main assumptions:

- During a fuel shortage that is expected to be temporary (weeks), services, businesses and individuals will try to maintain normal output despite fuel shortages
- Market behaviors will act to decrease fuel consumption by raising prices
- Prices rise faster than consumption decreases such that the cost of fuel (price times consumption) increases due to the disruption

The assumed relationship between fuel price and amount consumed is shown in **Figure 32**. In this figure, the horizontal axis is the fraction of normal consumption. That is, at the normal level of consumption the value on the horizontal axis is 1 and the price of fuel is $3.50 per gallon. When fuel consumption declines to 40% of normal, the price tops out at $10.00 per gallon. The relationship shown in **Figure 32** was informed by price data from an actual 2004 Phoenix fuel disruption.[27]

[27] http://www.doney.net/aroundaz/gas_lines.htm.

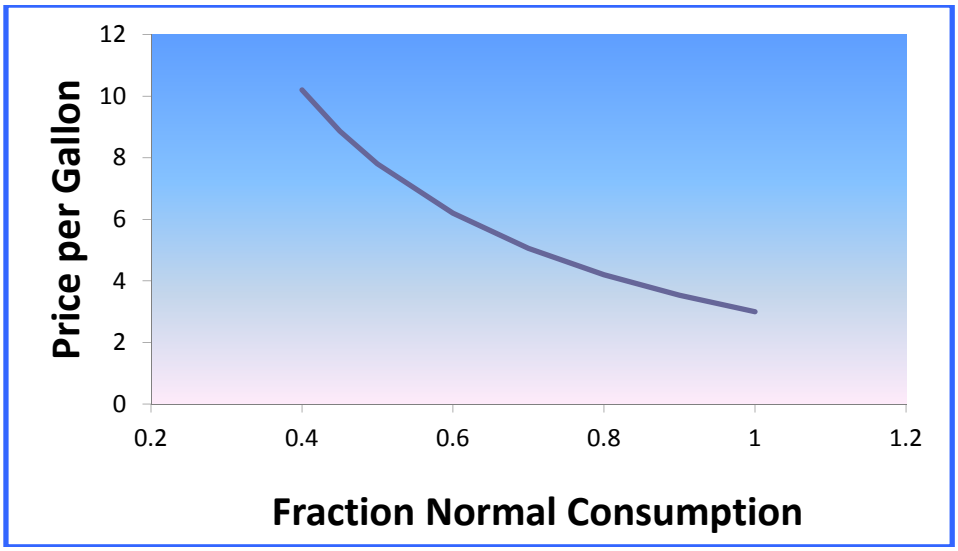

Figure 32. Assumed demand curve for fuel.

Given the price versus consumption curve shown in **Figure 32**, and the amount of fuel consumed from each distribution terminal on each day (as provided as output from the network model), it is easy to calculate the daily cost of fuel. The difference in this cost and the cost that would have occurred during an undisturbed period is the added fuel cost. The added fuel cost is the resilience metric calculated by the consequence model

D.3 Calculating the Consequence Metrics for the Scenario Earthquake

The three models were used to calculate a resilience metric for the scenario earthquake.

The first step is estimate a distribution of repair times for each of the four damaged pipelines. Such a distribution could be calculated by running the repair model multiple times, which each run using a set of plausible values for the model parameters. The multiple runs would represent the uncertainty in those parameters. For this example, we did not understand the distribution of those parameters well enough to provide a good example. Therefore we used an assumed distribution of repair times (**Figure 33**). This is a log-uniform distribution ranging from a repair time of one week to one year. This distribution is skewed toward shorter repair times because the horizontal axis is on a log scale.

We sampled this distribution to get 30 repair times and used those repair times as input to the network model. The result of these 30 simulations is a histogram (**Figure 34**) of the simulated number of barrels of fuel not consumed due to the disruption. **Figure 35** is therefore the performance indicator for this scenario. Note the horizontal axis of this histogram is also on a log scale, so the range of the intervals increases to the right.

Applying the consequence model to each of the 30 network model simulations results in the consequence metric is the likelihood of additional fuel costs. This metric is shown as **Figure 36**.

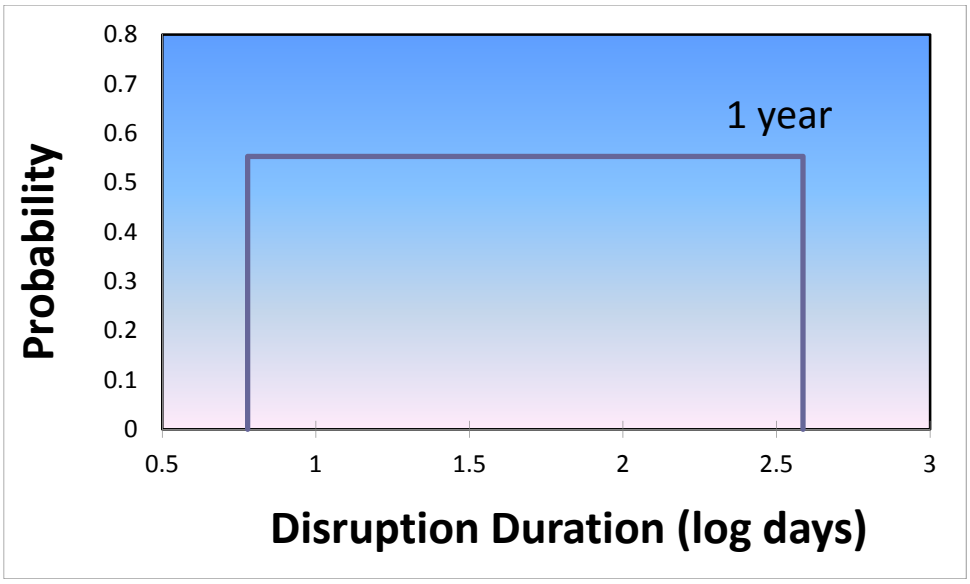

Figure 33. Probability distribution of repair times for each of for pipelines.

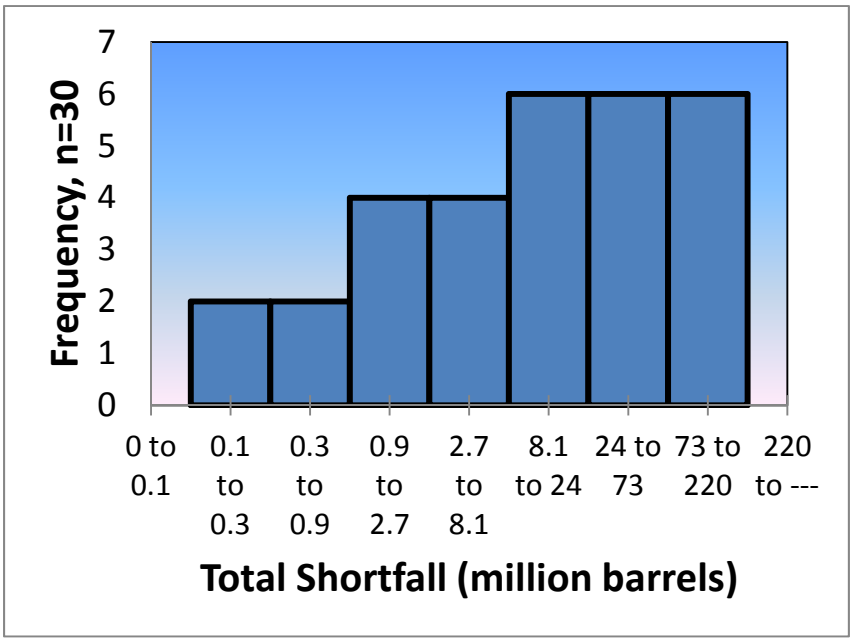

Figure 34. Histogram of 30 sampled repair times and resulting shortfall.

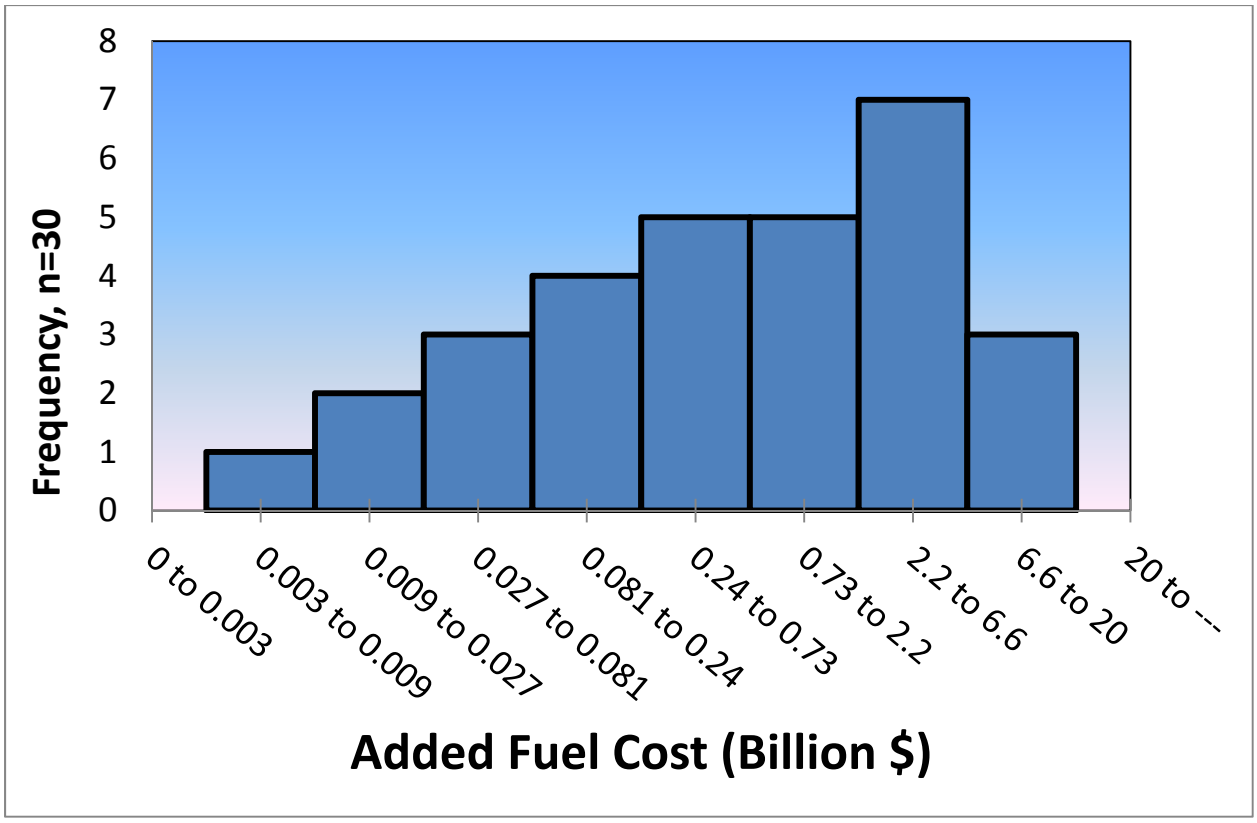

Figure 35. Histogram of 30 sampled repair times and resulting fuel cost increase.

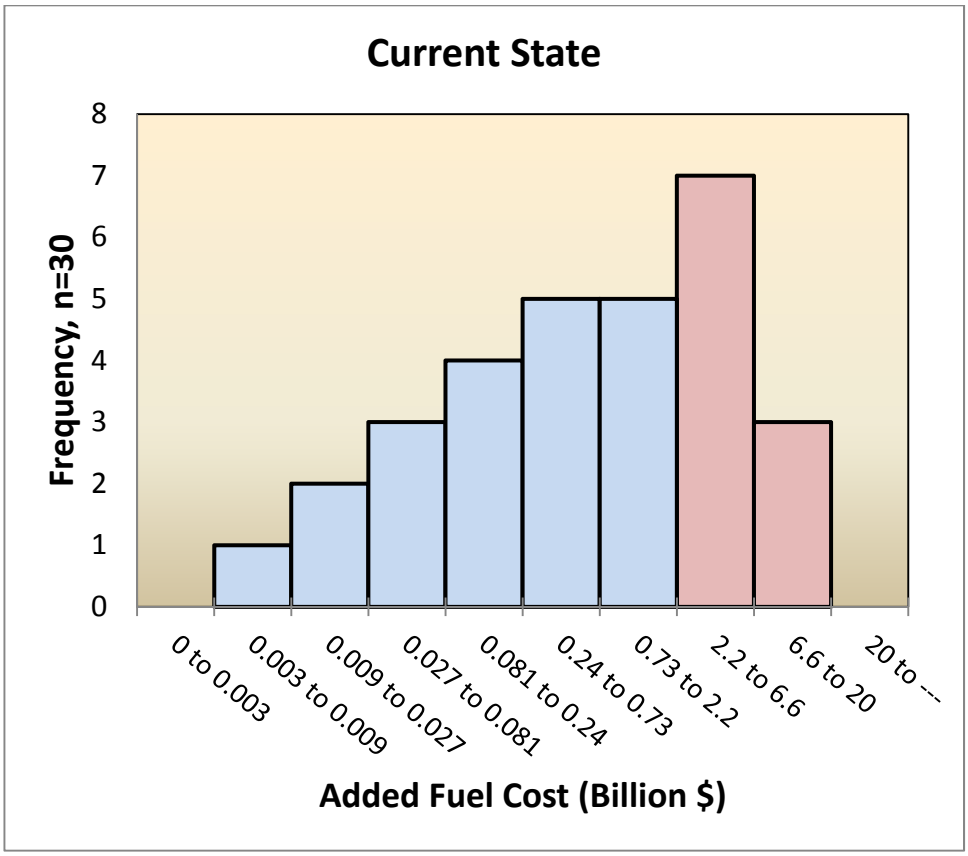

Figure 36. Resilience metric for this analysis, shown as a histogram.

D.4 Evaluating Investment to Increase Resilience

Two of four transmission pipelines are located in more favorable geologic conditions with respect to earthquake damage. Here we ask the question of how much the resilience of the petroleum infrastructure to the scenario earthquake could be increased by investing in re-engineering the TEPPCO and Midvalley pipelines such that they would suffer a down time of only one week for this earthquake. To answer this question, we repeat the calculation of the consequence metric using this assumption.

Figure 37 shows the re-calculated resilience. Comparing this result to **Figure 36** shows that the investment has reduced the likelihood of the added fuel cost of exceeding $2.3 Billion from 1/3 to 1/10.

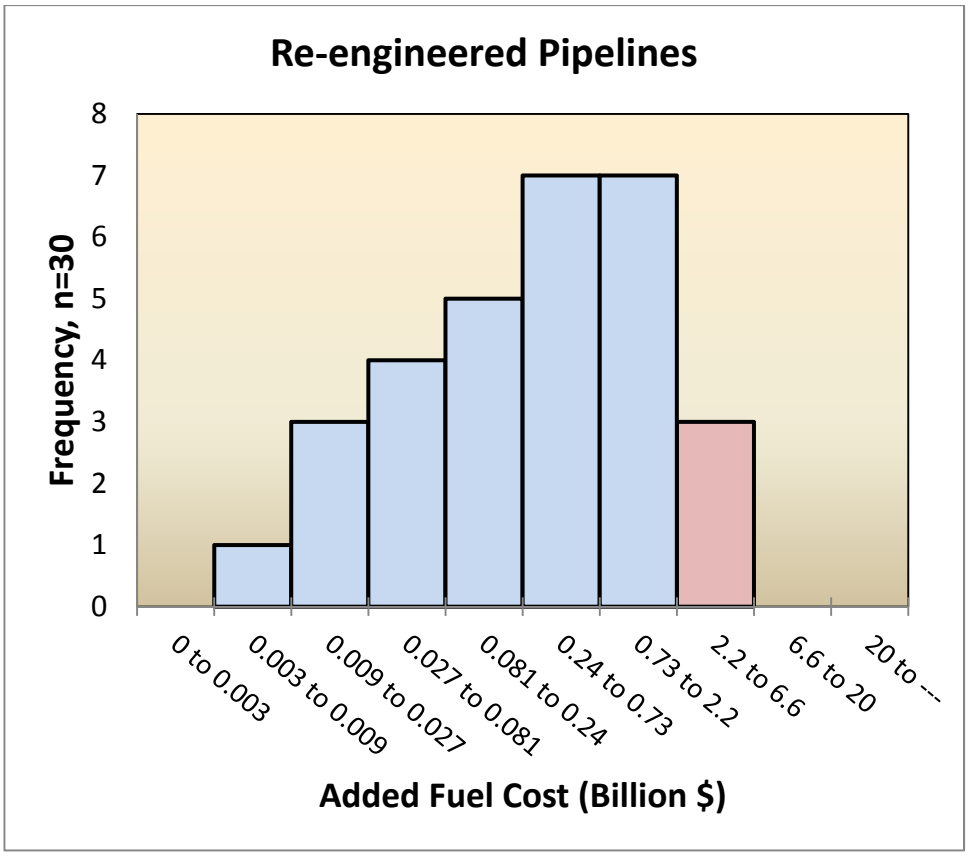

Figure 37. Improved resilience metric if two pipelines were relocated, shown as a histogram.

D.5 Conclusions/Recommendations

We suggest that a network model is the best way to estimate fuel shortages that result from disruptions to infrastructure components. In order for stakeholders to have confidence in such a model, it is necessary to reach agreement that:

- The model formulation represents system behavior well enough to calculate useful resilience metrics.
- The network definition (infrastructure capacities and connections) represents the actual state of the system well enough to calculate useful resilience metrics.
- For the petroleum use case, we calculated a single consequence metric (additional fuel cost). Additional research is required to (1) identify other metrics could be used to represent the impact of fuel shortages and (2) methodologies for populating these metrics.

Page intentionally left blank.

APPENDIX E: NATURAL GAS USE CASE

The natural gas use case presented in this section exemplifies how resilience metrics can be applied to the natural gas infrastructure. Here, we evaluate resilience by calculating the overall financial impact on the economy that natural gas delivery shortfalls would cause due to a natural disaster, a 7.8 magnitude earthquake at the San Andreas Fault near the Salton Sea. An engineering assessment for this type of earthquake was performed and results show that it would damage three important transportation corridors around the southern California area. **Figure 38** shows a map of the natural gas infrastructure around the San Andreas fault. The earthquake is assumed to take place sometime in December, correlating to high natural gas useage.

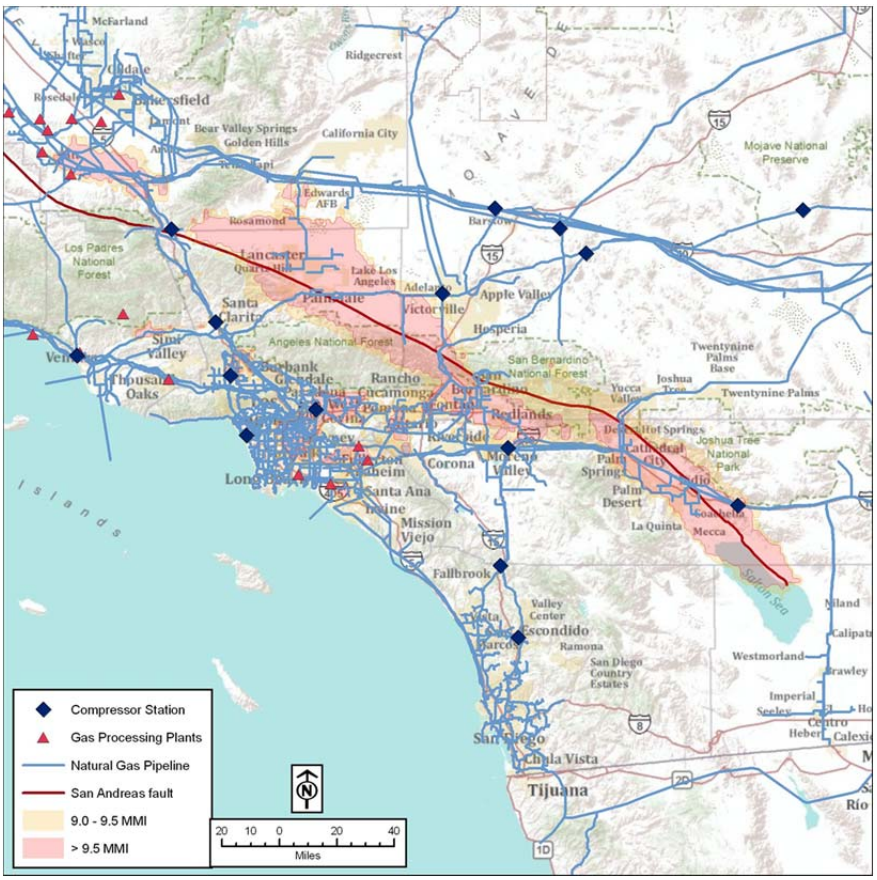

Figure 38. Map of the Southern California natural gas infrastructure in relationship to the San Andreas Fault.

The resilience of the Southern California basin is quantified under uncertain transportation corridor repair times and when the available natural gas storage facility is operated in two distinct ways. First, baseline resilience is quantified when storage withdrawals are restricted to historical values. Second, the operating policy changes by not restricting gas withdrawals from the storage facility. Resilience changes are assessed by comparing the economic impact for these two scenarios.

A network model of the North American natural gas infrastructure, the Gas Pipeline Competition Model (GPCM), is employed to calculate gas flows under normal and disrupted conditions. GPCM includes all major pipeline systems and uses market clearing in order to determine pipeline flows. This basic economic principle is also known as "competitive, partial equilibrium model" in economics literature. Its flow algorithm allows the network to adapt to disruptions. For instance, as price increases due to shortage demand is reduced and production is stimulated. GPCM also includes storage models and rerouting capabilities.

Network flows are calculated using GPCM for normal conditions, disrupted conditions with no storage withdrawal restrictions, and disrupted conditions with unrestricted storage withdrawals. These results for the Southern California area are summarized in **Figure 39**. Flow values correspond to millions of cubic feet of natural gas per day.

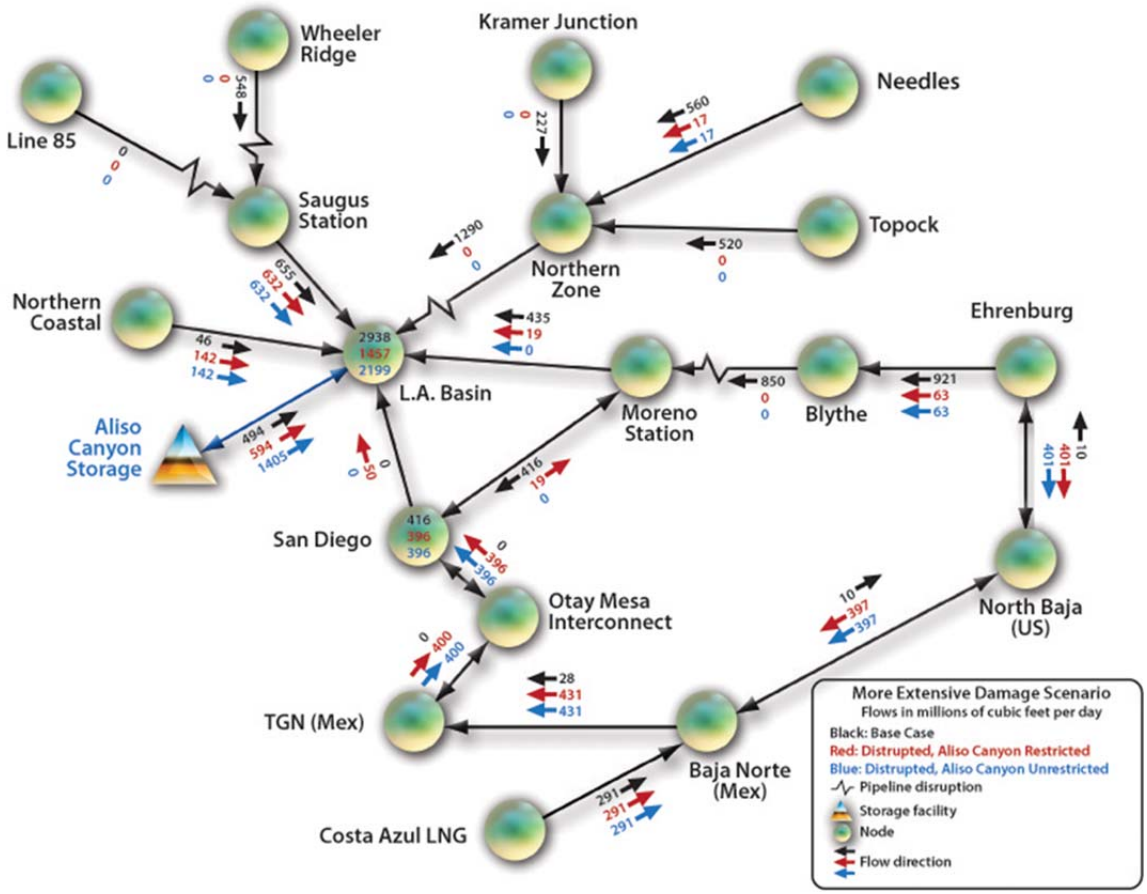

Figure 39. Southern California network flow results summary for normal conditions (black), disrupted conditions with storage withdrawal restrictions (red), and disrupted conditions with no storage withdrawal restrictions (blue).

The storage facility corresponds to the privately owned Aliso Canyon Storage. As natural gas prices increase due to the shortage caused by the three damaged transportation corridors, owners of the natural gas stored in this facility may wish to restrict the rate of withdrawal for economic gain or for any other reasons. This situation is modeled in the first case where withdrawals are

assumed to occur at rates seen in December of previous years under normal conditions. Under these conditions, model results show a natural gas supply approximately 50% below normal to the L.A. basin.

Many sources of uncertainty could be studied in this hypothetical natural-disaster scenario. One such source of uncertainty is the time that it would take for crews to repair the three natural gas transportation corridors. The repair time would be a function of the damage sustained by the pipelines such as number of breaks, access to damaged pipeline sections, availability of spare components, to mention a few. For illustration purposes, assume the uncertainty in repair times can be described using a normal distribution with a mean of 1 week and a standard deviation of 0.5 weeks. Repair costs are considered negligible in this analysis.

We proceed to calculate the shortage amount by subtracting the natural gas flow under normal conditions from the natural gas flow under disrupted conditions, which totals 1,481 MMcf/day. We assume this corresponds to the amount of additional natural gas needed in the L.A. basin that results in no economic impact when the earthquake occurs.

Next, we assess the effects of natural gas shortages on the Southern California economy. Again, for illustration purposes, assume that historical natural gas prices for different end uses under normal conditions in December provide a good proxy for the economic impact caused by the shortage. Historical prices for December 2013 were obtained from the energy information administration (EIA) website and are shown in Table 5.

Table 5. Natural Gas Prices by Sector for California in December 2013

Sector	Price ($/Mcf)
Residential	10.02
Commercial	8.27
Industrial	7.14
Transportation	4.41
Electric Generation	5.14

Additionally, information about monthly natural gas consumption per end use was obtained for the period between December 2013 and February 2014. We assume that usage per end sector of natural gas in California is maintained after the earthquake occurs.

We use Monte-Carlo simulation to estimate the economic impact by taking 1,000 samples from the repair time probability distribution described above. For each sample, the economic impact is calculated by aggregating the shortage per end use for the period of time that transportation corridors remain damaged. **Figure 40** shows the histogram of economic impact for the case described above. We refer to this result as “baseline resilience” for this use case. The expected economic impact is $325.1 million.

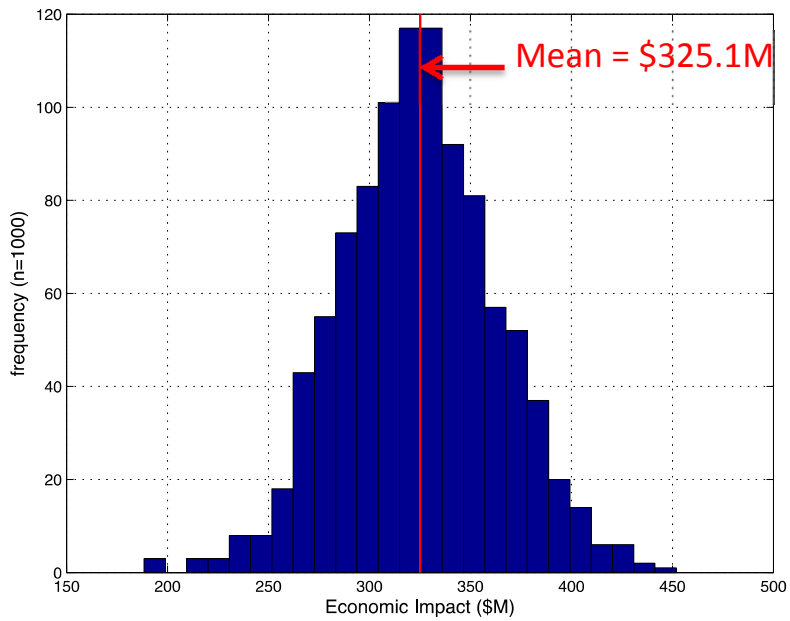

Figure 40. Histogram of economic impact shows baseline resilience of natural gas use case with a expected value of $325.1M.

The resilience metrics enable comparison between operating and planning strategies. In this natural gas use case we compare the baseline resilience when energy storage withdrawals are restricted with an alternative where storage withdrawals are unrestricted and only limited by the maximum flow rate of the storage facility physical components. Again, 1,000 samples from the repair time probability distribution are employed. Results are shown in **Figure 41**. The expected economic impact decreases to $163.1 million.

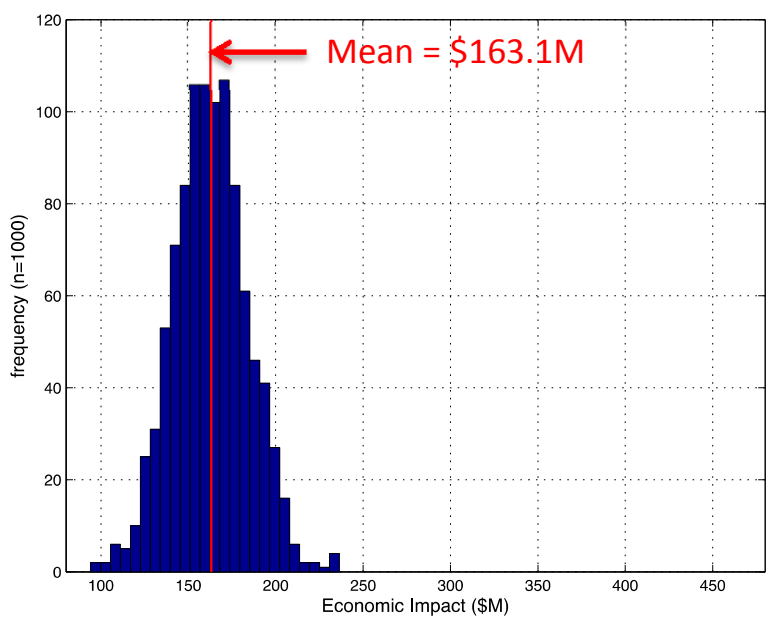

Figure 41. Histogram of economic impact with unrestricted natural gas withdrawals shows natural gas use case modified resilience

This example illustrates the use of resilience metrics to compare operating policies in the wake of a natural disaster that damages the natural gas infrastructure. Although much research is needed to refine several of the assumptions used in this analysis, we can conclude that operation of the Aliso Canyon storage facility has a major impact in the resilience to an earthquake in the San Andreas Fault.

R&D is needed to produce a more realistic resilience assessment. For instance, the model employed here is calibrated for use under normal conditions; uncertainty models for repair times and other sources of uncertainty are not readily available; no general consensus exists on how to translate from natural gas fuel shortages to resilience metrics such as economic impact.

Page intentionally left blank.

APPENDIX F: METHODS TO NORMALIZE RESILIENCE METRICS

Normalization is necessary in order to compare resilience values across cities or infrastructures. An example of such normalization would be to take lost commercial revenue and divide by the number of businesses affected (or total revenue, or capita, etc). The result, expected lost revenue per business could then be used to compare across electric infrastructure of different sizes and compositions.

As we have demonstrated in the use cases, resilience metrics can help guide operations and investments in order to improve energy systems' resilience. For instance, when comparing two different portfolios of infrastructure upgrades, resilience metrics permit a direct comparison of the resulting resilience levels and inform decision makers of which option results in greater resilience. However, resilience metrics should also permit comparison of an energy infrastructure's current state, including a quantitative understanding of that state with respect to other infrastructure so that a qualitative judgment can be made. For instance, the local council of a small town in a Midwest state is interested in knowing the current resilience of the electric grid in its city. The local power company performs studies following the framework presented in this work and brings back an answer to the local council. The answer is a set of numbers. This set contains resilience in terms of several different consequences such as lost commercial revenue ($), lives at risk and total recovery effort ($). These numbers correspond to the expected values of the probability distribution found by the power company for each of the consequences listed. These numbers by themselves do not offer a basis for a qualitative judgment (i.e., very resilient, not resilient) unless they can be compared with the electric infrastructure resilience from other towns or cities. Differences between infrastructure, population, load composition, etc. will inevitably result in different resilience values.

Page intentionally left blank.

APPENDIX G: COMBINING/CONSOLIDATING RESILIENCE METRICS FOR A SYSTEM

As defined by our resilience framework, the "value" of a resilience metric is a *distribution* of consequences (which can include restoration and recovery costs)—defined relative to a set of predefined disruption scenarios. In practice, it is very difficult to communicate and compare/contrast distributions to decision makers. Instead, summary statistics are used to transform the distribution into a scalar quantity, which can be easily interpreted. The most commonly used example of such a summary statistic is the mean of the distribution. However, alternative summary statistics may focus more on tail-oriented features of a distribution. An example of a range of summary statistics is given as follows:

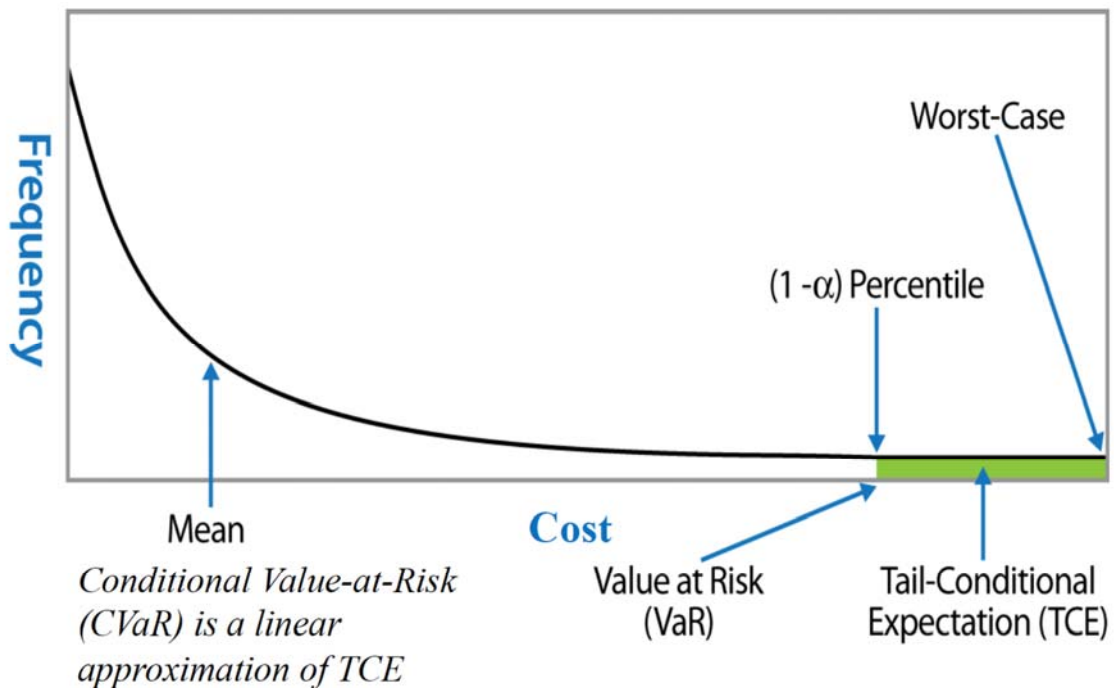

Figure 42. Waiting on caption.

In planning contexts, tail-oriented summary statistics are often preferred, to protect against "hundred-year flood" high-consequence, low-probability events. At a minimum, they are often combined in a linearly weighted manner with the expectation statistic. In the graphic above, the VaR and TCE statistics are defined in terms of a tail quantile of the distribution, e.g., 5%. VaR is then the cost of the 95% most costly disruption events. In contrast, TCE is simply the mean over the 5% most costly disruption events.

Understanding and communicating the relationship between different summary statistics is a major research challenge, involving both technical and social dimensions. For example, while different tail-oriented statistics are often correlated, the degree of correlation is problem-specific. Further, understanding the exact nature of the correlation often yields insights into the problem at

hand. Fundamentally, the problem resolves to one of multi-dimensional data analysis and visualization, which is an active research area.

Beyond collapsing consequence distributions to summary statistics, we must address the issue of analyzing distinct types of consequences simultaneously: for example, economic losses versus recovery costs. Given the disparate stakeholder preferences that are inherent in any comprehensive resilience analysis, some form of multi-objective performance analysis will ultimately be required. The most basic form of multi-objective analysis involves assigning weights to individual metrics (already transformed by a summary statistic), and computing the corresponding weighted sum. However, because this technique induces significant complications for interpretation of the resulting scalar quantities by decision makers, more general multi-objective analyses are preferred. These include, for example, generation and presentation of a "Pareto front", which shows the relationship between disparate metrics in the form of a multi-dimensional scatter-plot. For two and three dimensions, analysis of the trade-offs between different objectives is straightforward. However, interpretation of higher-dimensional trade-offs requires dimension reduction, which in turn complicates presentation and understanding of the metric relationships. Similar to the research challenges underlying summary statistic correlation analysis, this problem is fundamentally one of multi-dimensional data analysis and visualization.

DISTRIBUTION

4 Department of Energy, Office of Electricity Delivery and Energy Reliability
Attn: Dan Ton

4 Department of Energy, Office of Energy Policy and Systems Analysis
Attn: Kate Marks

1	MS1033	Charles Hanley	06112
1	MS1104	Carol Adkins	06100
1	MS1137	Robert Jeffers	06921
1	MS1137	Charles Rath	06921
1	MS1137	La Tonya Walker	06926
1	MS1138	Tom Corbett	06924
1	MS1140	Ross Guttromson	06113
1	MS1140	James Ellison	06113
1	MS1140	Cesar Silva-Monroy	06113
1	MS1188	Jared Gearhart	06131
1	MS1188	Katherine Jones	06131
1	MS1188	Dean Jones	06131
1	MS1362	Jean-Paul Watson	01464
1	MS9054	Robert Hwang	08300
1	MS0899	Technical Library	09536 (electronic copy)

Sandia National Laboratories