

QUANTIFYING IMPROBABILITY

**An Analysis of the Lloyd's of London
Business Blackout Cyber Attack Scenario**

National Security Report



Susan Lee | Michael Moskowitz | Jane Pinelis



JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

QUANTIFYING IMPROBABILITY

An Analysis of the Lloyd's of London *Business Blackout* Cyber Attack Scenario

Susan Lee

Michael Moskowitz

Jane Pinelis



JOHNS HOPKINS
APPLIED PHYSICS LABORATORY

Copyright © 2018 The Johns Hopkins University Applied Physics Laboratory LLC. All Rights Reserved.

Distribution Statement A: Approved for public release; distribution is unlimited.

Contents

Figures.....	v
Tables.....	vi
Summary.....	vii
Bulk Power System Background.....	1
Erebos Scenario Summary.....	3
Analysis Description.....	4
Analysis Approach.....	4
Scenario Interpretation.....	5
Data Set Generation.....	8
Simulation Implementation.....	9
The Erebos Scenario Analysis.....	9
The Random Selection Strategy.....	9
The Ordered Attack Strategy.....	11
Beyond Erebos.....	12
The Erebos Scenario Quantified.....	15
Analysis Limitations.....	17
Conclusions.....	17
Implications for Grid Defense.....	17
Implications for Quantitative Vulnerability Assessment.....	19
Implications for Policy.....	22
Appendix A Makeup of the Grid Segment Used in the Analysis.....	25
Appendix B List of Plants Used in the Analysis.....	31
Acknowledgments.....	41
About the Authors.....	41

Figures

Figure 1. Continental US Interconnections and NERC Regions.....	2
Figure 2. NERC Reliability Coordinators and Balancing Authorities	3
Figure 3. PCR for a Fixed Control Room Access Success Rate of 50% as a Function of All Three Factors	6
Figure 4. Convergence of Results	9
Figure 5. GAR for the Erebus Scenario vs. PCR—Fifty-Generator Limit	12
Figure 6. Comparison of Target Sets: Number of Grid Assets Needed for Goal Achievement— Fifty-Generator Limit	12
Figure 7. Number of Generators Damaged at Goal Achievement Using Random Selection— No Generator Limit, 18,000 Megawatt Capacity	13
Figure 8. Effect of a Generation Limit on GAR	14
Figure 9. Effect of a Generation Limit on Number of Grid Assets Needed for Goal Achievement	14
Figure A-1. Number of Generators per Plant in NPCC and RFC (Excluding 48 in Mountain View and 73 in Edison Sault)	25
Figure A-2. Number of Generators by Capacity in NPCC and RFC	26
Figure A-3. Number of Generators by Type in NPCC and RFC	26
Figure A-4. Generator Capacity by Type	27
Figure A-5. Total Plant Capacities in Descending Order (Nuclear, Solar, and Wind Excluded)	28
Figure A-6. Plant Capacity vs. Number of Generators in the Plant.....	28
Figure A-7. Unequal Distribution of Generator Capacity within Plants	29

Figure credits:

Information in Figures 1 and 2 from North American Electric Reliability Corporation (NERC), <http://www.nerc.com/Pages/default.aspx>.

Tables

Table 1. Words of Estimative Probability.....	7
Table 2. Study Parameters.....	10
Table 3. Goal Achievement for Random Selection Strategy—Fifty-Generator Limit	10
Table 4. Goal Achievement Using an Ordered Attack Strategy for the Erebos Scenario.....	11
Table 5. Goal Achievement for Random Selection Strategy—No Generator Limit	13
Table 6. Goal Achievement for the Ordered Selection Strategy—No Generator Limit	13
Table 7. Low-PCR Adversaries Targeting Any Size Plant in NPCC and RFC.....	15
Table B-1. Plants with Generators with at Least 100 Megawatts of Capacity.....	31

Table credits:

Table 1 modified from Sherman Kent, *Words of Estimative Probability* (Washington, DC: Studies in Intelligence, Central Intelligence Agency, 1964).

Summary

The 2015 Lloyd's of London and the University of Cambridge Centre for Risk Studies report, *Business Blackout*, hypothesized a cyber attack on the US electric grid, and estimated economic consequences between \$60 billion and \$200 billion.¹ While the hypothesized consequences were severe, the risk entailed cannot be calculated without an assessment of the probability that such an attack could occur. *Business Blackout* offered a qualitative assessment, based on subject matter expertise, that the attack was possible but improbable. We used the cyber attack scenario described in *Business Blackout* to demonstrate how a probabilistic assessment could be used to quantify the likelihood that the scenario could occur.

The analysis showed that the cyber attack scenario, interpreted as best as possible from the description in the report, was both possible and improbable. Going further to assess the sensitivity of that result to some of the parameters of the scenario, we discovered that the scenario was more restrictive than it seemed. By relaxing some arbitrary constraints on the hypothesized adversaries, or by giving the adversaries slightly more skill or risk tolerance, the attack became considerably more probable. We found that an assessment of vulnerability to a grid-wide attack, like the cascading blackout hypothesized in *Business Blackout*, must consider the grid as a system, since the results were significantly impacted by the actual distribution of grid assets with particular characteristics within the region of study. We made a number of conservative assumptions, favoring the adversaries, to make our results reflect the worst case. For example, we did not consider the adversaries' risk of detection during their campaign.

A nonlinear response to parameter changes implied that small changes in the level of defense could have an outsized impact on the probability that the adversaries would achieve their goal. A defense did not need to be perfect to essentially preclude the adversaries from achieving their goal. Making a relatively small number of changes (relative to requiring changes grid-wide) could have a big impact, if we understand where to make them.

The analysis is subject to the limitations inherent in probabilistic risk assessment. Despite its limitations, the analysis fulfilled its intended purpose of demonstrating how a probabilistic risk assessment could be used to quantify the likelihood that a cyber attack will achieve a specific effect. Better data would be needed to make results from such an analysis definitive. Some of this could be collected today; we will need a bold and imaginative approach to generate the required data. Finally, if grid defense is to be considered holistically, as a property of the grid rather than individual components, we conclude that very significant changes will be needed to both determine and enforce a policy.

¹ S. Ruffle, E. Leverett, A. Coburn, J. Copic, S. Kelly, T. Evan, D. Ralph, M. Tuveson, O. Bochmann, L. Pryor, and J. Z. Yeo, *Business Blackout: The Insurance Implications of a Cyber Attack on the US Power Grid* (Cambridge, UK: Lloyd's of London and the University of Cambridge Centre for Risk Studies, 2015).

Many reports on the possible effects of a large-scale cyber attack on the electric grid exist.¹ They consistently predict very significant disruptions in national security, the economy, and quality of life. Most of these reports include a qualitative assessment of the vulnerability of the grid to such an attack. These assessments consist largely of references to the known small-scale incidents involving individual grid assets, or even incidents involving other physical plants using industrial control systems (ICSs). While known incidents speak to the possibility of cyber attack, they do not shed light on the risk of a large-scale attack. Cyber risk is a function of threat (an actor with intent), consequences, and probability that an attack of the required caliber can be executed. Quantifying the probability of the success of large-scale cyber attack is hard, due to lack of precedent and the changing nature of threats and vulnerabilities; however, without a quantitative assessment of the probability of occurrence, the risk posed by even very significant consequences cannot be determined.

In 2015, Lloyd's of London and the University of Cambridge Centre for Risk Studies published an influential report, *Business Blackout*, on the economic consequences of a hypothesized cyber attack on the US electric grid.² The resulting outage was assumed to be extensive, and in some cases, long-lived. The economic consequence was estimated to be in the range of \$60 billion to \$200 billion. The *Business Blackout* report did not address the probability of an attack like the one it hypothesized; its stated purpose

was to draw out the implications for the insurance industry if such an attack did occur. Nevertheless, the report included a plausible attack scenario, supported by subject matter experts and some analysis (not specified) to ensure that technical aspects of the scenario were possible. The report deemed the attack scenario to be improbable, but not impossible.

The purpose of this analysis is to derive a slightly more quantitative assessment of the risk posed by the scenario in the *Business Blackout* report. Rather than asking whether this scenario is possible, this analysis addresses the likelihood that the scenario could happen as described. Since the parameters of the attack were somewhat ambiguous and arbitrary, the analysis is extended to study the sensitivity of the outcome to the specifics of the hypothesized attack and to explore the effect of improving cyber defenses. Ultimately, the work presented here illustrates an approach to understanding the vulnerability of the grid to large-scale cyber attack.

Bulk Power System Background

The bulk power system (BPS) comprises generation and transmission assets that provide electricity to distribution points from which electricity is then delivered to individual consumers. In the continental United States, the BPS is divided into three large interconnections: the Eastern Interconnection, the Western Interconnection, and the ERCOT Interconnection, as shown in Figure 1. The primary feature of an interconnection is that it operates synchronously; that is, electric power is generated and transmitted at the same frequency, phase, and voltage, within small bounds. To maintain synchronism, the demand for electricity (load) and the generation of electricity must be balanced (equal) at all times. Load fluctuates a certain amount constantly (e.g., every time a light is turned on, load increases slightly). These normal fluctuations are compensated with an increase (or decrease) in generation; however, there is a limit to the size and rapidity of generation response. A sudden, large deviation from a balanced state anywhere in

¹ Stuart Madnick, "Preparing for the Cyberattack That Will Knock out the U.S. Power Grids," *Harvard Business Review*, May 10, 2017; Robert K. Knake, *A Cyberattack on the U.S. Power Grid*, Contingency Planning Memorandum No. 31 (New York: Council on Foreign Relations Center for Preventive Action, April 2017); and Ted Koppel, *Lights Out: A Cyber Attack, A Nation Unprepared, Surviving the Aftermath* (New York: Crown Publishers, 2015).

² S. Ruffle, E. Leverett, A. Coburn, J. Copic, S. Kelly, T. Evan, D. Ralph, M. Tuveson, O. Bochmann, L. Pryor, and J. Z. Yeo, *Business Blackout: The Insurance Implications of a Cyber Attack on the US Power Grid* (Cambridge, UK: Lloyd's of London and the University of Cambridge Centre for Risk Studies, 2015).

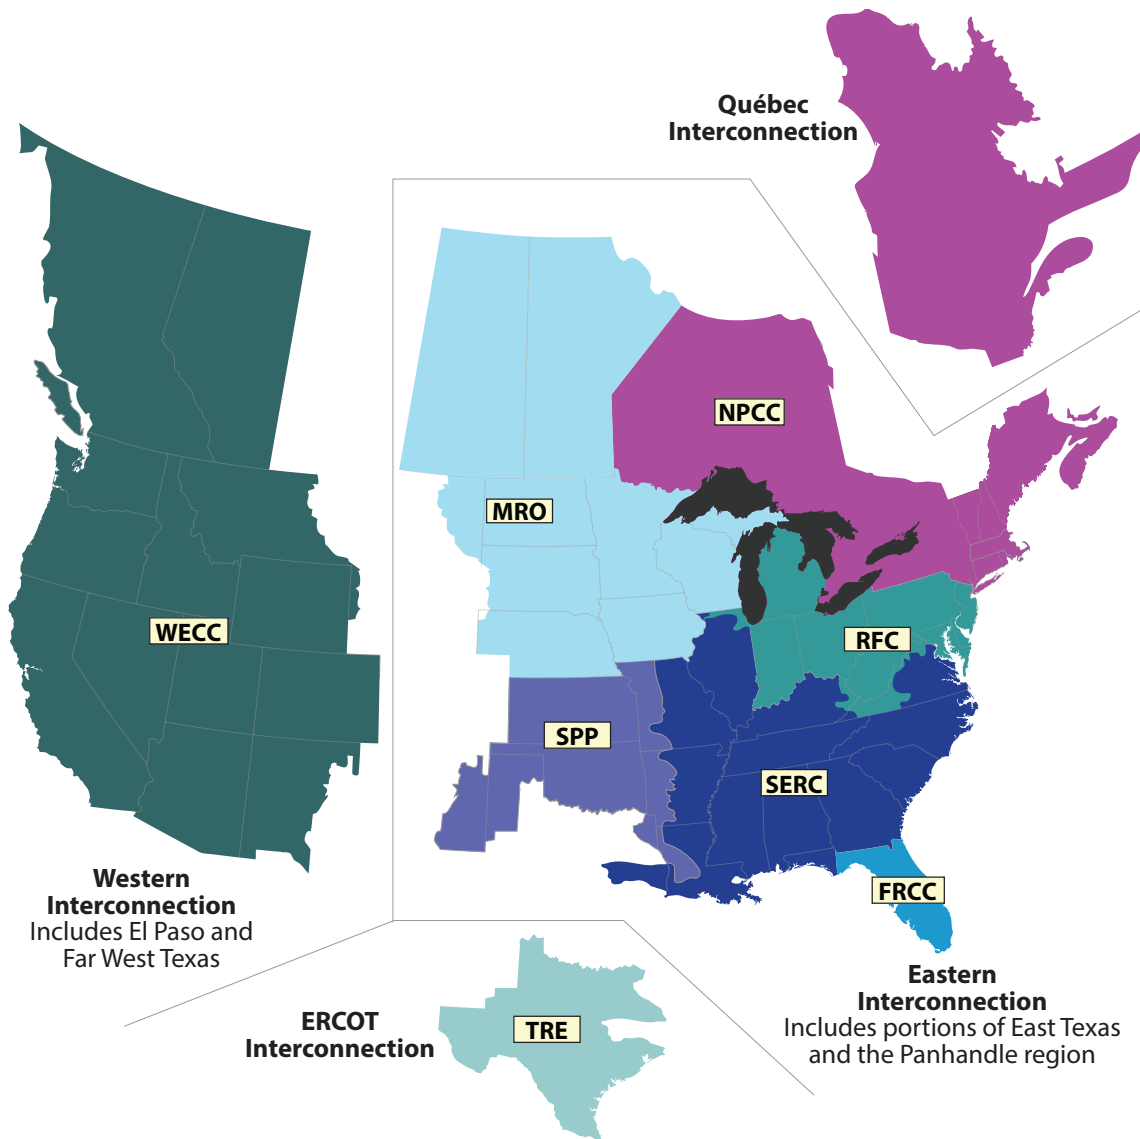


Figure 1. Continental US Interconnections and NERC Regions

the interconnection will propagate across the entire interconnection. As it propagates, limits set on the operating parameters of assets may be violated. As relays designed to protect the assets from damage from out-of-limit conditions trip, the disturbance in balance can become more severe. Although this is rare, large portions of an interconnection have been shut down as a result of a local imbalance—for example, the August 2003 Northeast blackout.³

³ *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations* (U.S.-Canada Power System Outage Task Force, April 2004).

The grid is further subdivided into several administrative and operational domains. The North American Electric Reliability Corporation (NERC) divides the grid into eight regions (also depicted in Figure 1), two of which, Northeast Power Coordinating Council (NPCC) and the ReliabilityFirst Corporation (RFC), figure in the *Business Blackout* report. NERC regions are responsible for maintaining NERC reliability standards within their boundaries.⁴ Day-to-day operations of the BPS are the responsibility of two

⁴ “Key Players,” North American Electric Reliability Corporation, <http://www.nerc.com/AboutNERC/keyplayers/Pages/default.aspx>.

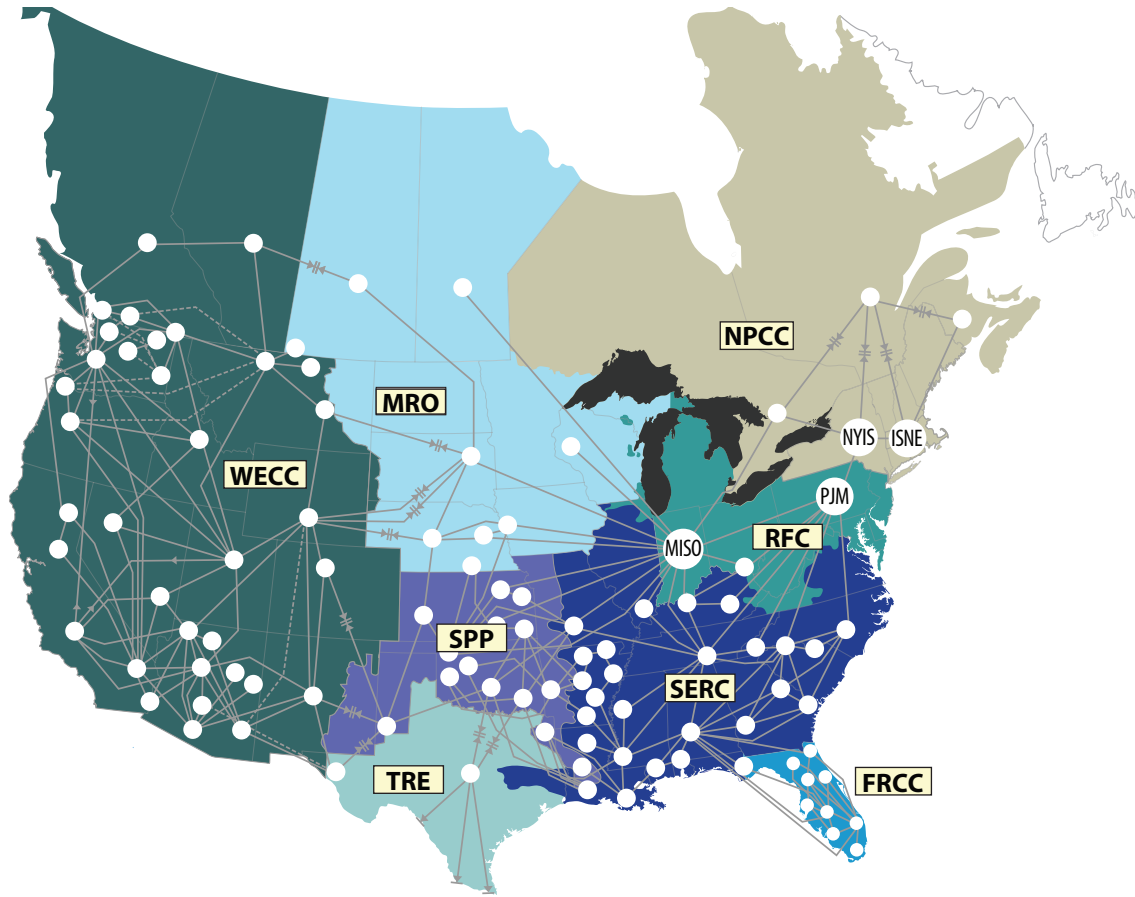


Figure 2. NERC Reliability Coordinators and Balancing Authorities

other entities defined by NERC: the reliability coordinator and the balancing authority (see Figure 2). The reliability coordinator is the highest operating authority and is responsible for managing exchanges among the balancing areas (BAs) in its own area and negotiating exchanges among BAs in neighboring reliability coordination areas. A balancing authority is responsible for real-time operation of assets within its own BA to meet its own demand, as well as keeping exchanges with neighboring BAs within the negotiated limits. These roles can be taken on by separate organizations or a single organization. The entities discussed in the *Business Blackout* report and used in this analysis, Pennsylvania-Jersey-Maryland (PJM), New York Independent System Operator (NYIS), Midcontinent Independent System Operator (MISO), and Independent System Operator New England (ISNE), perform the functions of both

roles. The boundaries of NERC regions and the two operational control areas are not necessarily aligned.

Erebus Scenario Summary

The *Business Blackout* report described the Erebus Scenario based on input from cyber and grid subject matter experts. An unidentified group with ample resources and contacts within the hacking community wishes to significantly disrupt the US economy. It chooses to use a cyber attack on two NERC regions: NPCC and RFC. The objective of the cyber attack is to damage enough generation simultaneously to initiate a cascading failure that blacks out the New York/DC corridor; damage, rather than simple shutdown, is chosen to slow recovery from the blackout.

The scenario description states that the area under attack has an hourly peak load of 194,000 megawatts

based on July 2014 US Energy Information Administration (EIA) data. In the scenario, the attack must damage enough generators to remove about 18,000 megawatts, approximately 10 percent of the generating capacity in NPCC and RFC, to initiate the cascade.⁵ The Erebos adversaries attempt to compromise control rooms in power plants to gain control over at least 18,000 megawatts of generation, taking advantage of the fact that some control rooms manage multiple generators. In some control rooms that were successfully compromised, the adversaries could access vulnerable generators and overcome the protective devices that would prevent the damaging attack.

In the scenario, after several months of activity, the adversaries were prepared to attack seventy vulnerable generators. Only fifty of them were successfully damaged at the time of the attack, but the capacity of these fifty generators totaled 18,000 megawatts. The sudden loss of approximately 10 percent of the generating capacity in the target area had the desired effect of blacking out both New York City and Washington, DC. The blackout extended over a region encompassing some or all of fifteen states. No mention is made of any impact in Canada.

Analysis Description

The following sections describe our probabilistic risk assessment approach, and how we cast the Erebos Scenario within that framework. We also carefully describe how we chose the data set of plants and generators used in the study to best approximate the set used in the *Business Blackout* report.

Analysis Approach

As the *Business Blackout* report makes clear, a cyber attack of this nature would be very challenging

for attackers to carry out successfully.⁶ Achieving the ultimate goal (a large cascading blackout) is a function of attacker skill, defender skill, and, in part, luck. For example, malware is very sensitive to the exact configuration of the target host; an exploit that works in one place may fail in another, seemingly similar, place. The particular operations occurring immediately prior to the attack could cause the attack to fail for reasons entirely unrelated to the malware itself. There is no necessary relationship between malware effectiveness and the size of the generation facility. This analysis treats the outcome of the Erebos Scenario as a random variable, dependent on the success of the adversaries' attacks on individual plants and the interaction of their malware with individual generators.

Attacker and defender skill is represented by a single uniform distribution. This single parameter controls the adversaries' success at implanting malware that gives them the ability to damage generators. In this analysis, attacker skill was additionally represented as a strategy for picking plants to attack. To represent the possibility of unanticipated circumstances at the time of the attack, a random subset of successfully implanted malware fails to damage the targeted generator.

The underlying distribution of the size of generators and generators among plants strongly influences the results. For this analysis, we approximately reproduced the actual topology of the NPCC and RFC regions studied in *Business Blackout* (see Appendix A). We tested the effect of using different subsets of NPCC and RFC facilities as the target set; however, results for another region and grid makeup will not be the same.

Theoretically, all combinations of fifty generators summing to 18,000 megawatts of capacity in the target area could be computed; however, the combinatorics are large and computationally

⁵ Ruffle et al., "Annex B: The US Electricity Grid and Cyber Risk to Critical Infrastructure," *Business Blackout*.

⁶ Ruffle et al., "Annex C: Constructing the Scenario – Threats and Vulnerabilities," *Business Blackout*.

expensive. We used a stochastic simulation to replicate the conditions described in the Erebos Scenario and estimate the probability of a successful attack. The simulation is described in the Simulation Implementation section of this report.

Scenario Interpretation

For the purposes of the economic analysis in *Business Blackout*, the details of the cyber attack—the number of plants that are attacked, the number of control rooms successfully compromised, the number of vulnerable generators, the distribution of vulnerable generators among control rooms, etc.—are less important than the hypothesized impact on the delivery of electricity. The introduction to the report states that the description of the cyber attack is purposefully, and understandably, vague; however, the scenario includes certain details that cyber and industry subject matter experts agree are plausible.

The scenario description states that (1) 10 percent of control room compromises yielded access to vulnerable generators; (2) one hundred control rooms were compromised, but in 57 percent of these, generators were protected against attack; and (3) seventy vulnerable generators were infected.⁷ Elsewhere⁸ the report states that seventy plants would probably have to be compromised to control enough capacity to cause the hypothesized effect. The scenario description goes on to say that of seventy infected generators, fifty are successfully damaged on the day of the grid attack. In a table describing the outage caused by the attack, fifty damaged generators are equated with 10 percent of all vulnerable generators,⁹ and in a computation of property loss, the fifty damaged generators are said to be 7 percent

of the total number of generators.¹⁰ These descriptors are impossible to reconcile with each other, and with the actual makeup of the grid in the NPCC and RFC regions. To perform this analysis, some interpretation was required.

Four separate factors implicit in the Erebos Scenario description impact the probability that the adversaries can achieve their goal:

- (1) Success in gaining access to the plant's control room
- (2) Given control room access, success in gaining access to vulnerable generators
- (3) Given access to vulnerable generators, success in overcoming protective mechanisms
- (4) Given a vulnerable, unprotected generator, successful execution of the damaging attack at the moment of the adversaries' choosing

The first factor, access to a plant control room, is strictly a function of defender and attacker skill. Some control rooms are better defended than others. There are likely no control rooms that are perfectly defended against all adversaries. Even if all technical defenses are perfect, there remains social engineering or the insider. The capability adversaries bring to the campaign will determine the level of defense that can be overcome.

In the Erebos Scenario, the adversaries aim to damage generators. Like any mechanism, generators are susceptible to damage if they are operated improperly. Generators of different types are not necessarily subject to the same type of damage. Because they are expensive and have a long lead time for acquisition, generators are carefully protected against improper operation; they have devices that sense unsafe conditions and put the generator into a safe condition (e.g., disconnect from the grid) before damage can ensue. To damage a generator, the adversaries must first disable the protective devices

⁷ Ruffle et al., *Business Blackout*, 11.

⁸ Ruffle et al., "Annex C: Constructing the Scenario – Threats and Vulnerabilities," *Business Blackout*.

⁹ Ruffle et al., *Business Blackout*, 15.

¹⁰ Ruffle et al., *Business Blackout*, 30.

and then instigate an unsafe operation specific to the type of generator.

The adversaries in the *Business Blackout* scenario use some Aurora-like attack to damage generators.¹¹ Aurora exploits an inherent vulnerability of any large rotating mechanical device, such as a turbine generator, to a sudden large resistance to its rotation; connecting a turbine generator to the grid out of synchronism (at a different voltage, frequency, or phase) is like slamming on a very large brake. In our analysis, we removed solar-voltaic generators (no rotating parts) and wind-powered turbines (not directly connected to the grid) from the target set. Most turbine generators directly connected to the grid are likely vulnerable to damage from an Aurora-like attack. Since we do not know the specifics of each generator's design, we assumed all were vulnerable.

In the past, protection from damage was provided by electromechanical devices not subject to cyber attack. Now, nearly all protective devices are small special-purpose computers that accept commands (some of them, remote commands) to control their operation. If the adversaries are properly positioned, these devices can be overcome.

In our analysis, we use a single factor, the Plant Compromise Rate (PCR), to represent factors 1–3: probability of compromising the control room, probability of finding vulnerable generators, and probability of overcoming protection against damage. We use PCR to determine the adversaries' ability to compromise a control room, and we assume that all generators are vulnerable and inadequately protected. Since the Erebus Scenario description states that 10 percent of control room compromises yielded access to vulnerable generators, we assume that a PCR of 10 percent is the equivalent to adversary success assumed in *Business Blackout*.

To cyber subject matter experts, a 10 percent probability of compromising control rooms may seem low. If PCR represented only the probability of compromising a control room, it might be low. In our analysis, PCR includes two additional factors. Our results are representative of any combination of the three factors that together multiply out to the PCR value. Even relatively high success rates for individual factors can result in a low PCR. Figure 3 shows that adversaries successful at compromising control rooms 50 percent of the time would have an overall PCR near or below 10 percent if either or both of the other two factors is 50 percent or less.

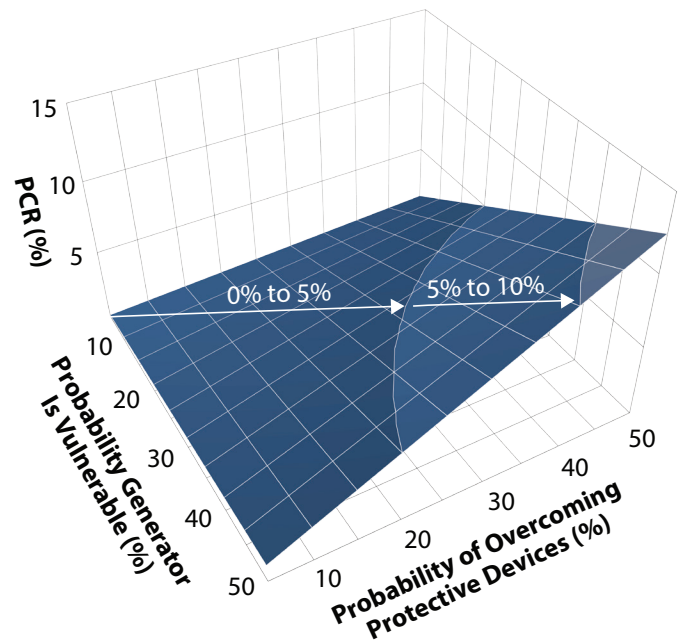


Figure 3. PCR for a Fixed Control Room Access Success Rate of 50% as a Function of All Three Factors

The fourth factor—successful execution of the grid attack at the moment of the adversaries' choosing—takes into account all the uncertainties in malware effectiveness described in the Analysis Approach section of this report. The malware effectiveness was unambiguous in the Erebus grid attack description; that is, out of seventy generators infected, fifty were actually damaged as a result of the attack. We used the Erebus ratio, 5/7, in our analysis.

¹¹ Mark Zeller, "Myth or Reality – Does the Aurora Vulnerability Pose a Risk to My Generator?" *Proceedings of the 64th Annual Conference for Protective Relay Engineers* (New York: IEEE, 2011), 130–136.

As noted in the Erebos Scenario description, the longer the adversaries work to increase the number of generators under control, the greater their chances of being discovered or wiped out by some unfortunate (from the adversaries’ perspective) change in configuration at a plant where they had previously succeeded. The analysis presented in this report ignores the time aspect and assumes that once the adversaries are in control, they remain in control. We use the number of successful attacks on plants and generators as an indication of adversary activity that could lead to detection, but we do not assign a likelihood of detection. The assumptions that the adversaries both remain in control of all previously infected assets and successfully elude detection are best-case scenarios for the adversaries.

The selection of a target set is a fifth, highly impactful, factor in the scenario. Although *Business Blackout* never explicitly stated that generator vulnerability was coincident with generator capacity, there were numerous explicit and implicit references to generators of 100 megawatts or greater capacity. Further, any set of fifty generators summing to a combined capacity of 18,000 megawatts would have to contain many large generators. We assumed the adversaries only attacked plants that contained generators with capacities of 100 megawatts or greater. Although control rooms often control generators of widely varying sizes, we assume the adversaries attempt to damage only the large (at least 100 megawatt) generators. Because the strict interpretation of the Erebos Scenario lets the adversaries damage only fifty generators, this is the best-case assumption for the attackers.

In summary, we chose to interpret the Erebos Scenario as follows:

- The adversaries attack plants with turbine generators with a capacity of 100 megawatts or greater.
- The adversaries succeed in compromising some percentage of the plants they attack (PCR).
- All generators in a compromised plant are susceptible to the attack and are vulnerable to adversary control.
- The adversaries attempt to control only generators with at least 100 megawatts.
- The adversaries maintain control over a generator once they have attained it.
- A fraction (5/7) of all the susceptible adversary-controlled generators are damaged when the grid attack commences.
- If the damaged generators have a total capacity of 18,000 megawatts or more, the adversaries achieve their goal and start a cascading blackout.

The analysis will assess the adversaries’ Goal Achievement Rate (GAR) while varying the PCR, the adversaries’ attack strategy, the target set, and the number of generators the adversaries control when they launch the overall grid attack. It is the adversaries’ goal achievement—that is, creating widespread damage and starting a cascading blackout—that the *Business Blackout* report deemed improbable. To compare a quantitative GAR to a qualitative judgment, we looked to Sherman Kent’s classic treatment of *Words of Estimative Probability* (see Table 1), with minor modifications.¹²

Table 1. Words of Estimative Probability

Level of Certainty	Definition
Certain	>99%
Almost certain	93%, give or take about 6%
Probable	75%, give or take about 12%
About even	50%, give or take about 10%
Probably not	30%, give or take about 10%
Improbable	7%, give or take about 5%
Impossible	<1%

¹² Sherman Kent, *Words of Estimative Probability* (Washington, DC: Studies in Intelligence, Central Intelligence Agency, 1964).

We note that the probability of an event may be low, but if the number of opportunities is high, even low-probability events can happen. For this study, we assume that once the adversaries initiate their overall attack on the grid, they will not have another opportunity for many years. Essentially, we assume that the adversaries can make only one attempt to launch the overall grid attack and that they succeed with the GAR we estimate.

Data Set Generation

Since our analysis aims to quantify the probability that the Erebus Scenario could occur, we attempted to recreate the data set of companies, plants, and generators used by the *Business Blackout* analysts. The *Business Blackout* analysts used data from the EIA to perform their analysis. We also used EIA detailed operating data from 2014 to obtain a data set that approximately replicates these data from *Business Blackout*.¹³

The *Business Blackout* report described the attack area as two NERC regions: NPCC and RFC. NPCC and RFC comprise all or parts of seven NERC BAs. The regions cover some or all of twenty-two states, as well as Ontario, Quebec, and New Brunswick in Canada. As described in the Bulk Power System Background section of this report, NERC regions are not necessarily coincident with reliability coordination regions or BAs,¹⁴ entities that are more closely tied to the operation of specific parts of the power grid. The description of the attack area as NERC regions made it somewhat difficult to determine which generation facilities the authors of the *Business Blackout* report considered in the creation of their scenario.

The *Business Blackout* report states that, within the target NPCC and RFC regions, there are 150

companies owning 261 power plants that control 676 generators with capacities greater than 100 megawatts. It further reports that the July 2014 hourly peak load was 194,000 megawatts.

The EIA website provides Excel spreadsheets listing US utilities (companies) and plants cross-referenced to NERC region and BA. From this, the analysts generated a list of plants located within the US portions of the NPCC and RFC regions. For the purposes of this analysis, the list was pruned to plants within the PJM, NYIS, MISO, and ISNE BAs. Only five plants in the EIA database were associated with the NPCC and RFC regions but were outside these four BAs.

The EIA website also provides a spreadsheet of generators cross-referenced to plants, but not to NERC region and BA. The analysts created a list of generators in the NPCC and RFC regions by cross-referencing with the plant list described in the preceding paragraph. These two EIA spreadsheets were not entirely consistent. Approximately 10 percent (288) of the plants in the NPCC and RFC plant spreadsheet did not appear in the generator spreadsheet. This analysis uses only the plants for which both region and generator data were available, yielding a list of power plants in the US portions of NPCC and RFC and the number and individual generator capacities at each.

The resulting list was filtered to remove nuclear generating facilities. *Business Blackout* explicitly excluded nuclear plants. As previously mentioned, we also excluded wind and solar generation because those technologies would require a different type of damaging attack.

We created a subset of this list containing only plants with one or more 100-megawatt generators, based on nameplate capacity. Since the realized summer and winter capacities are lower than the nameplate capacity, this is a worst-case assumption. To compare with the data from *Business Blackout*, this subset comprised 157 companies and 244 power plants controlling 698 generators with at least 100-megawatt

¹³ Available in Excel spreadsheet format at “Form EIA-860 Detailed Data with Previous Form Data (EIA-860A/860B),” US Energy Information Administration, <https://www.eia.gov/electricity/data/eia860/>.

¹⁴ NERC, “Key Players.”

capacity. This is close, but not identical, to the figures cited in *Business Blackout*. The discrepancy may be explained by the possible use of different choices to account for EIA data set discrepancies, the use of summer versus nameplate capacities, or the exclusion of solar and wind technologies. At any rate, the inclusion of more large generators controlled by fewer plants is a worst-case assumption for risk analysis. The list of power plants used in this analysis is in Appendix B.

Simulation Implementation

Using the R programming language, we developed a function that would allow us to simulate scenarios where the adversaries attempt to control plants and damage generators. The target set of generators and plants, the attackers' strategy (random or ordered selection), the PCR, and the malware effectiveness ratio are all arguments passed to the function, allowing for exploration of a number of different scenarios.

The function generates each individual realization of the scenario by first ordering the target plant list according to the attackers' strategy. For random selection, the target set was randomly ordered. For the ordered strategy, the target set was always listed in order from largest to smallest plant by plant capacity.

Each plant in the target set was assigned a random number in the interval $[0,1]$. If the random number was less than the PCR (expressed as a decimal), we count the plant as one controlled by the adversaries for that scenario realization. The generators with more than 100 megawatts within each controlled plant were assigned another random number on the interval $[0,1]$. If that random number is less than the malware effectiveness ratio ($5/7$ or 0.714), we count the generator as damaged. Starting with the first plant in the target set, the program calculates a running total of the metrics of interest: the total capacity of the damaged generators, the number of unique plants controlled, and the total number of damaged generators. For each realization, the running total is

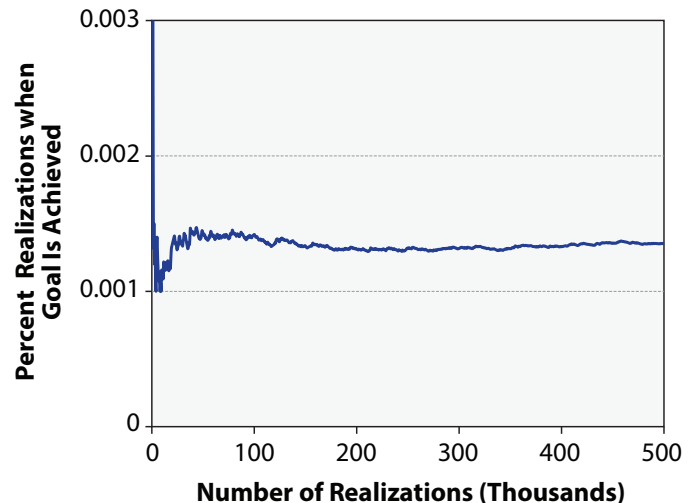


Figure 4. Convergence of Results

saved for the points at which the fiftieth generator is damaged, the points at which the target capacity of 18,000 megawatts is reached, and for the entire set. The metrics from each realization are aggregated for 500,000 realizations, to provide the results reported throughout the Erebos Scenario Analysis section in this report. Convergence analysis indicates that 500,000 realizations are sufficient for a stable estimate of the parameters of interest (see Figure 4).

The Erebos Scenario Analysis

In our analysis, we studied the adversaries' GAR under varying assumptions for strategy, number and size of generators, targeted plants, and PCR. In all cases, the attackers' goal was to damage 18,000 megawatts of generation capacity. The analysis results for each combination are described in the subsequent sections. Table 2 lists the combinations of parameters we studied and the section where the result can be found.

The Random Selection Strategy

The *Business Blackout* report did not indicate that the adversaries attacked plants in any particular order. We analyzed the success of a random selection strategy—that is, randomly selecting plants from the target set

Table 2. Study Parameters

Section	Strategy	Generator Number and Size	Plants Targeted	PCR (%)
The Random Selection Strategy	Plants attacked at random	≤ 50 generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts	100
The Random Selection Strategy	Plants attacked at random	≤ 50 generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts; largest 20 plants removed	100
The Ordered Attack Strategy	Plants attacked in order by total capacity	≤ 50 generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts	5–100
The Ordered Attack Strategy	Plants attacked in order by total capacity	≤ 50 generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts; largest 20 plants removed	5–100
Impact of Removing Generator Limit	Plants attacked at random	Any number of generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts	100
Impact of Removing Generator Limit	Plants attacked in order by total capacity	Any number of generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts	5–100
Impact of Removing Generator Limit	Plants attacked at random	Any number of generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts; largest 20 plants removed	100
Impact of Removing Generator Limit	Plants attacked in order by total capacity	Any number of generators; only generators ≥ 100 megawatts	Plants with generators ≥ 100 megawatts; largest 20 plants removed	5–100
Impact of Removing Target Set Limit	Plants attacked at random	Any number of generators; generators of any size	All plants of any size	100
Impact of Removing Target Set Limit	Plants attacked at random	Any number of generators; generators of any size	All plants of any size; largest 20 plants removed	100
Impact of Removing Target Set Limit	Plants attacked in order by total capacity	Any number of generators; generators of any size	All plants of any size	5 and 10
Impact of Removing Target Set Limit	Plants attacked in order by total capacity	Any number of generators; generators of any size	All plants of any size; largest 20 plants removed	5 and 10

until the Erebus Scenario is replicated. This could model attackers who send out a blanket set of phishing emails and then simply work on the first, random, set of “bites.” For this plant selection strategy, we gave the attackers 100 percent probability of success in compromising any given plant; the element of chance here is in the plants that are selected for attack.

In our early exploration of the data, we saw that when the adversaries achieved their goal, one or more of a small set of plants always appeared in the list of compromised plants. Not surprisingly, these were among the plants with the largest capacity. To understand the importance of these plants to grid defense, we analyzed how likely the adversaries were to achieve their goal if those plants were successfully defended. We created a variant of the target set by

removing the twenty plants with the highest total capacity. This is equivalent to giving the adversaries a 0 percent chance of compromising those plants.

The results are shown in Table 3. Even with a 100 percent success rate in compromising plants, the best GAR is about 7 percent. Removing the largest twenty plants reduced the GAR to less than 1 percent;

Table 3. Goal Achievement for Random Selection Strategy—Fifty-Generator Limit

Goal and Target Set	Statistics	All Target Plants	Top 20 Removed
50 generators; ≥18,000 megawatts total capacity	% goal achieved	6.8	0.3
	Average no. of plants	22.9	26.2
	Average no. of generators	47.2	48.4

this validates our observation that these plants figure heavily in success within the Erebos fifty-generator constraint. The average number of generators damaged when the goal was achieved was close to fifty, showing that the adversaries seldom succeeded without nearly fifty generators. These results indicate that the Erebos Scenario is both possible and improbable, as assumed in *Business Blackout*, at least if the adversaries use a random selection strategy.

The Ordered Attack Strategy

An element of attacker skill is strategy. To determine whether a better strategy than random selection could significantly increase the adversaries’ chances of achieving their goal, we assumed that the adversaries would attack plants in order from largest to smallest. The results for the ordered strategy are presented in Table 4 and Figure 5.

The difference between the random and ordered strategies is stark. When adversaries with a 100 percent PCR used the ordered strategy, they were 100 percent successful in achieving the goal, as opposed to using random selection where they achieved the goal, at best, 6.8 percent of the time. For the full target set, only the 5/7 successful malware execution ratio prevented the adversaries from achieving the goal in the theoretical minimum of seven plants. When the largest twenty plants were removed from the target

set, the adversaries with the 100 percent PCR needed to compromise 50 percent more plants to achieve the goal.

To assess the impact of plant compromise success rate on achieving the ultimate goal of damaging fifty generators with at least 18,000 megawatts of capacity, we also analyzed PCRs of 5, 10, 20, 40, and 80 percent. We believe that the 10 percent PCR represents the Erebos adversaries. At this PCR, even using an ordered selection strategy, the adversaries’ goal was achieved only 7 percent of the time; that is, success was possible but improbable, as described in *Business Blackout*. At a PCR of 20 percent, the adversaries’ GAR of 59.5 percent falls into the range of about even. At PCRs between 40 and 100 percent, the adversaries were nearly 100 percent successful at achieving the goal. At these high PCRs, the goal was met with an average number of generators far fewer than fifty. This indicates that the challenge is less in controlling a large number of generators than in controlling the largest generators.

Removing the largest twenty plants from the target set lowered the probability of goal achievement significantly. The GAR at 10 percent became negligible, and even at a 20 percent PCR, it was only about 7 percent. At the highest PCRs, 80 and 100 percent, the adversaries had still achieved the goal nearly 100 percent of the time. At all PCRs, the adversaries had to compromise more plants and

Table 4. Goal Achievement Using an Ordered Attack Strategy for the Erebos Scenario

Goal and Target Set	Statistics	PCR (%)					
		5	10	20	40	80	100
50 generators > 18,000 megawatts total capacity; all target plants	% goal achieved	0.05	7.0	59.5	98.5	99.9	100.0
	Average no. of plants	16.1	17.1	14.8	11.9	9.9	9.4
	Average no. of generators	42.8	44.3	41.6	36.4	34.7	30.6
50 generators > 18,000 megawatts total capacity; top 20 removed	% goal achieved	0.0002	0.1	7.3	54.7	99.0	99.99
	Average no. of plants	19.0	21.0	19.5	16.8	15.4	15.0
	Average no. of generators	44.0	47.3	46.8	44.1	38.8	36.9

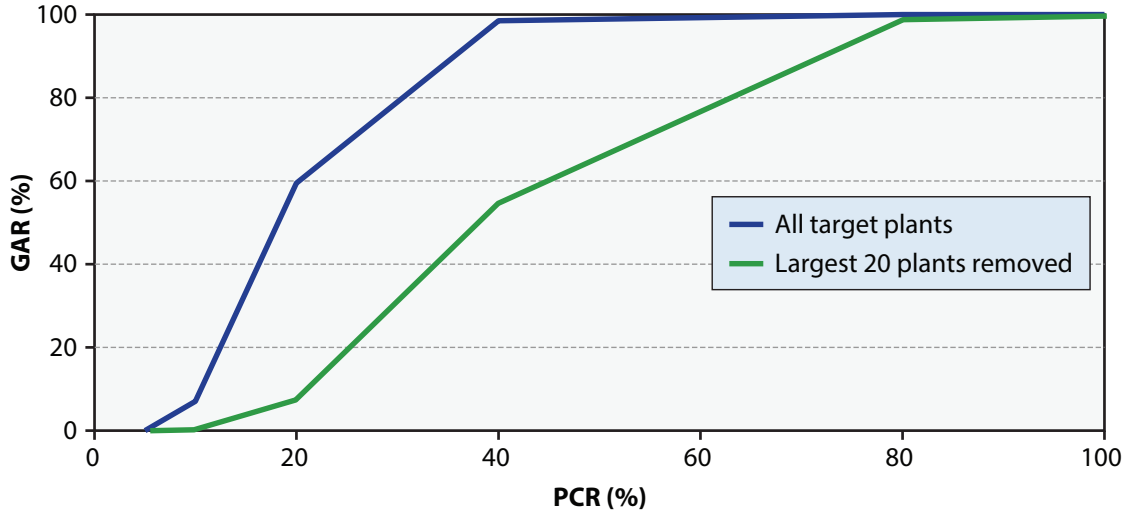


Figure 5. GAR for the Erebos Scenario vs. PCR—Fifty-Generator Limit

damage more generators to achieve the goal. As shown in Figure 6, this effect was more pronounced at the higher PCRs. When all plants were in the target set, the adversaries with 80 and 100 percent PCRs could achieve the goal within the first few large plants. When these plants were removed, the adversaries were forced to compromise more smaller plants. At lower PCRs, the adversaries often failed to compromise the largest plants anyway, and removing

those plants from the target set had less impact on the number of plants required for success.

Beyond Erebos

The Random Selection Strategy and Ordered Attack Strategy sections analyze the probability that the exact Erebos Scenario occurs; that is, the adversaries damage fifty generators accounting for a total of

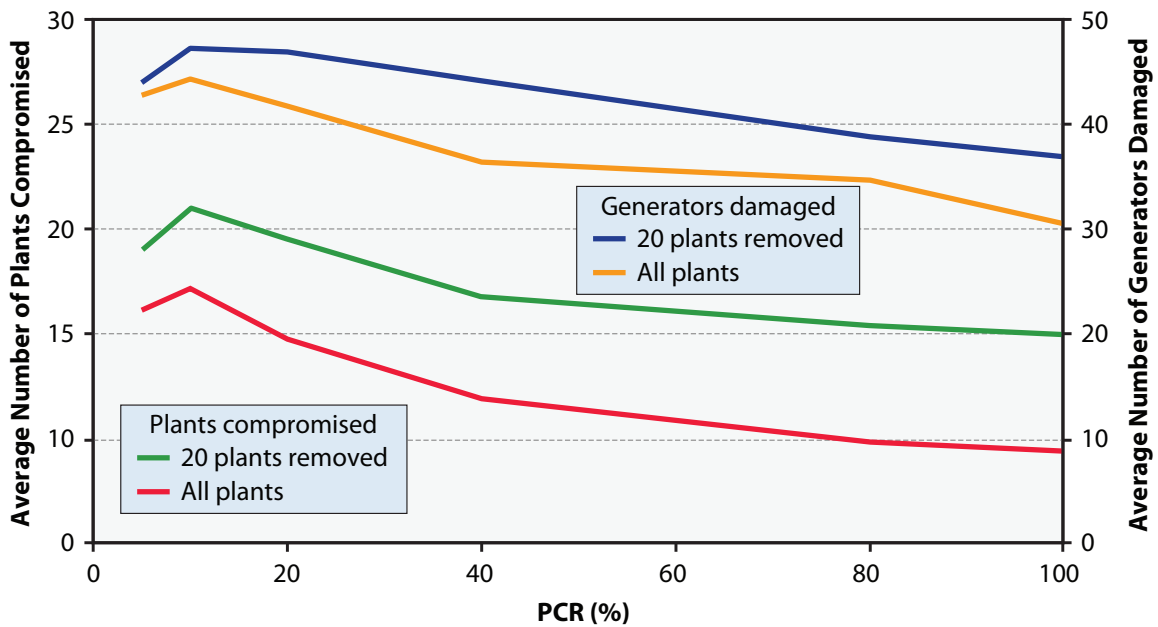


Figure 6. Comparison of Target Sets: Number of Grid Assets Needed for Goal Achievement—Fifty-Generator Limit

Table 5. Goal Achievement for Random Selection Strategy—No Generator Limit

Goal	Target Set	All Target Plants	Top 20 Removed
Any number of generators; $\geq 18,000$ megawatts total capacity	% goal achieved	100.0	100.0
	Average no. of plants	28.2	33.4
	Average no. of generators	62.3	70.0

18,000 megawatts or more in capacity. The Erebos adversaries stopped compromising control rooms and initiated their overall grid attack because they worried that continued activity would lead to discovery. Although the decision to stop at seventy generators controlled, and to ultimately damage fifty, was grounded in expert opinion, it was an arbitrary choice. In this section, we analyze the probability that bolder adversaries would achieve their goal.

Impact of Removing Generator Limit

We explored what happens if the adversaries continue to compromise plants and damage generators in the target set until they reach 18,000 megawatts of capacity or have attacked every plant in the target set and compromised as many as they can (given the PCR) without reaching the goal capacity.

The results for the random selection strategy with this new goal are shown in Table 5. Because the PCR is 100 percent, the adversaries are 100 percent successful in eventually damaging enough generators to take 18,000 megawatts out of the grid. The

average number of generators it took was considerably higher than fifty. Removing the largest twenty plants from the target set did not impact the GAR, but it further raised the number of plants and generators needed to reach the goal. The many failed attempts to reach the goal when limited to fifty generators (see the Random Selection Strategy section) versus 100 percent success with sixty to seventy generators shows that, *taken as a whole*, there are relatively few sets of fifty generators within the target set that total to 18,000 megawatts. As shown in Figure 7, there are a fairly large number of sets of sixty to seventy generators that total to 18,000 megawatts total capacity.

The results for the ordered strategy are shown in Table 6. The impact of strategy can be seen by

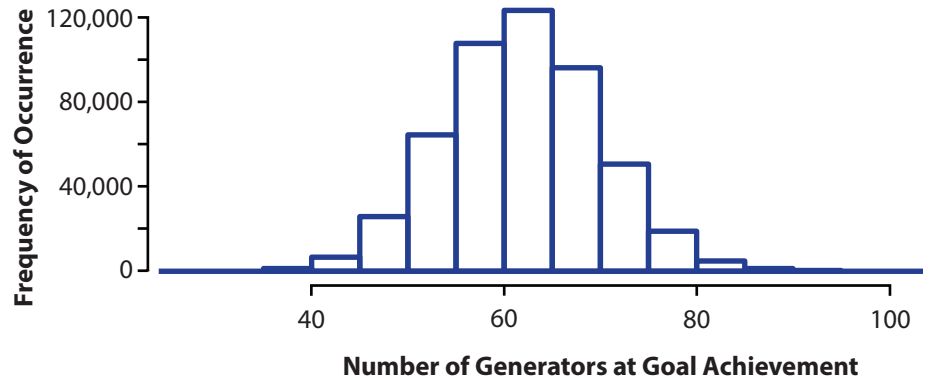


Figure 7. Number of Generators Damaged at Goal Achievement Using Random Selection—No Generator Limit, 18,000 Megawatt Capacity

Table 6. Goal Achievement for the Ordered Selection Strategy—No Generator Limit

Goal and Target Set	Statistics	PCR (%)					
		5	10	20	40	80	100
>18,000 megawatts total capacity; all target plants	% goal achieved	0.06	17.4	99.0	100.0	100.0	100.0
	Average no. of plants	16.8	19.9	16.9	12.0	9.9	9.4
	Average no. of generators	46.1	52.5	48.1	36.6	34.7	30.6
>18,000 megawatts total capacity; top 20 removed	% goal achieved	0.0006	1.3	85.8	100.0	100.0	100.0
	Average no. of plants	20.0	24.7	24.7	17.9	15.4	15.0
	Average no. of generators	50.3	60.7	61.8	49.9	39.0	36.9

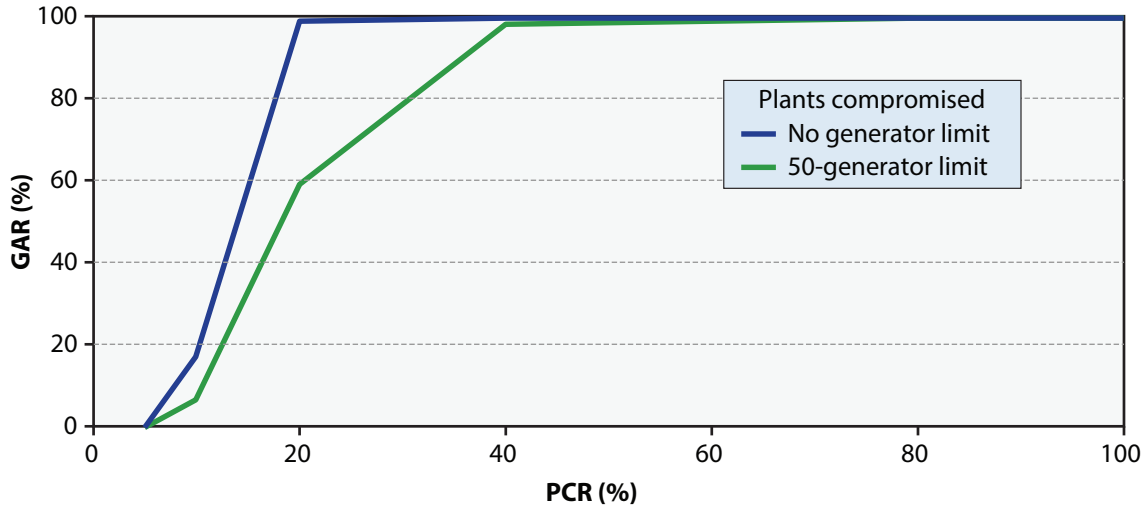


Figure 8. Effect of a Generation Limit on GAR

comparing the average number of grid assets needed for 100 percent GAR in Table 5 and Table 6. By working from the top down, adversaries cut the number of plants they need to compromise by a factor of three.

Figure 8 compares the GAR at various PCRs for the Erebos fifty-generator limit and for no generator limit. For the higher PCRs, where number of generators was never a limiting factor, the removal of this constraint had no impact. At the lowest PCRs,

the adversaries were able to increase their success rate. At the nominal Erebos PCR of 10 percent, the adversaries achieved the goal about 17 percent of the time, versus about 7 percent when limited to fifty generators. Still, using our *Words of Estimative Probability* (see Table 1), a GAR of 17 percent falls midway between “improbable” and “probably not.” Figure 9 shows the additional grid assets that the adversaries needed to make these advances. Again, removing the largest twenty plants from the target set

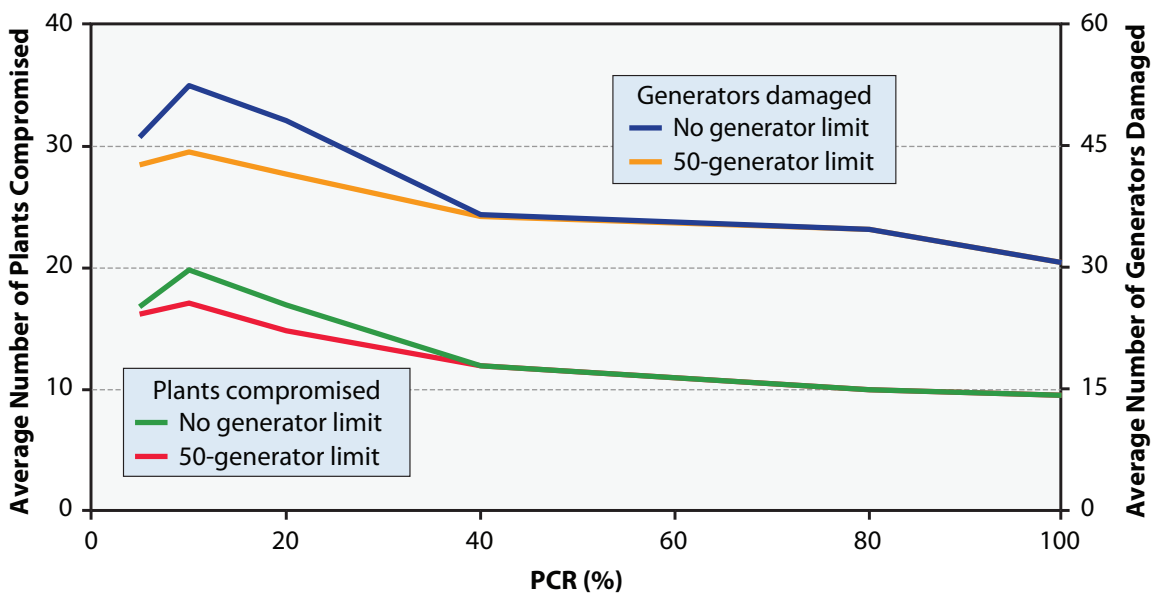


Figure 9. Effect of a Generation Limit on Number of Grid Assets Needed for Goal Achievement

suppressed goal achievement at the lower PCRs and increased the plant assets required for success.

Impact of Removing Target Set Limit

The target set used in the preceding sections (plants with generators at least 100 megawatts and only generators at least 100 megawatts) is not a stringent limitation, since it contains about 202,000 megawatts of the 270,000 megawatts of total capacity within the studied regions. Nonetheless, adversaries with 5 and 10 percent plant compromise success rates were unable to get a high GAR within this target set. This failure may be a function of sheer numbers: at a 5 percent PCR, the adversaries will compromise, on average, only about twelve plants out of the total 244 in the target set. While it is possible to achieve the goal while compromising only twelve plants, they must all be plants at the higher end of the range of total capacity. At a 10 percent PCR, the average number of plants compromised is twenty-four; even at twenty-four plants, a set with 18,000 megawatts of total capacity must include some of the larger plants.

To assess the success of an ordered selection strategy unconstrained by the particular set of plants, we repeated the analysis in the Ordered Attack Strategy and Impact of Removing Generator Limit sections with a target set that contained all the plants and all the generators (not just the 100-megawatt generators) in NPCC and RFC. This set includes 1,595 plants and 5,196 generators. The results are shown in Table 7. If limited to the Erebos Scenario fifty-generator set, the adversaries did no better than they did when constrained to the set of plants with generators larger than 100 megawatts. This is to be expected, since so much of the generation capacity is within that set. On the other hand, when allowed to accumulate generation until achieving the goal of 18,000 megawatts, the adversaries with a 10 percent PCR were able to succeed nearly 60 percent of the time. To do so, they had to compromise an average of 32 plants and damage 111 generators. If the Erebos adversaries were right to be

Table 7. Low-PCR Adversaries Targeting Any Size Plant in NPCC and RFC

Goal and Target Set	Statistics	PCR (%)	
		5	10
50 generators; > 18,000 megawatts total capacity; all NPCC and RFC	% goal achieved	0.01	1.02
	Average no. of plants	14.0	13.8
	Average no. of generators	45.2	45.7
Any number of generators; > 18,000 megawatts total capacity; all NPCC and RFC	% goal achieved	0.6	59.4
	Average no. of plants	28.5	32.7
	Average no. of generators	99.4	111.3
50 generators; > 18,000 megawatts total capacity; all minus top 20	% goal achieved	0	0.004
	Average no. of plants	—	18.0
	Average no. of generators	—	47.6
Any number of generators; > 18,000 megawatts total capacity; all minus top 20	% goal achieved	0.01	21.8
	Average no. of plants	34.3	46.6
	Average no. of generators	112.7	147.9

concerned about being discovered, this amount of activity could lead to detection and removal before the goal was reached.

Removing the largest twenty plants had the expected effect of reducing GAR and raising the number of plants compromised and generators damaged. In the 10 percent PCR case, the adversaries were only successful 22 percent of the time and had to compromise more than forty-six plants and damage 148 generators on average when achieving their goal.

The Erebos Scenario Quantified

The authors of *Business Blackout* stated that at least some sets of fifty generators totaling to more than 18,000 megawatts of capacity exist in the NPCC and RFC target region. We used random sampling to estimate the number. If adversaries with a 100 percent PCR randomly attack all the plants

containing generators with at least 100 megawatts in US portions of NPCC and RFC, they will come up with a set of fifty generators with 18,000 megawatts total capacity about 7 percent of the time. This percentage fits the definition of improbable offered in the Scenario Interpretation section of this report. On the other hand, the same adversaries allowed to sample randomly until they accumulate the requisite 18,000 megawatts of capacity will succeed 100 percent of the time with an average of only sixty-two generators. Just a tweak on the Erebos Scenario (sixty-two versus fifty generators) makes the accumulation of enough generation capacity to start the hypothesized cascading blackout about 50 percent, even without a strategy for targeting the largest plants.

Adversaries that are both skilled in cyber attack *and* select their targets randomly are an unlikely combination. We considered what happens when adversaries systematically attack the largest plants first, with varying levels of success. When the conditions of the Erebos Scenario were adhered to—that is, 10 percent successful compromise of plants with vulnerable generators, fifty generators damaged—organized adversaries will still succeed only about 7 percent of the time, again improbable. Allowed to damage as many generators within our target set as needed to get the requisite capacity, adversaries with a 10 percent PCR succeeded 17 percent of the time, still in the range of improbable. This presupposes a threat (adversary) willing to put in the effort for such a low probability of success.

At high PCRs, even when constrained to the fifty-generator limit of the Erebos Scenario, the adversaries achieved the goal nearly 100 percent of the time. When allowed to accumulate as many generators as needed to reach the needed capacity, even adversaries with a 20 percent PCR were successful nearly 100 percent of the time. For adversaries with somewhat better skills than the Erebos adversaries, the Erebos Scenario is probable, perhaps even certain. For our analysis, at somewhere

between a 10 percent and a 20 percent PCR, there is a tipping point for GARs. The nonlinearity of this result is interesting. If future, better-informed analysis confirms this phenomenon, there are very significant implications for grid defense.

Allowed to attack across the entire region, unfettered by artificial limits on number and size of generators, the Erebos adversaries with a 10 percent PCR will succeed about 60 percent of the time, between even and probable odds in *Words of Estimative Probability*. This probability of success could be attractive to an adversary. On the other hand, to achieve this success, the adversaries need an average of about one hundred generators. Our analysis explicitly ignored the element of time and the probability of detection that real adversaries would encounter with continued activity over a lengthy period. Further analysis, including this effect, could show that the attack is self-limiting, due to increased detectability, given the amount of activity required to control enough generation.

Because the attack in the Erebos Scenario was a damaging attack, it is difficult to perform. A low PCR may be most representative. On the other hand, if the goal were merely initiating a cascading blackout, the bar is lower; possibly all that would be required is access to an operator console in a control room to disconnect generators from the grid. The overall *risk* from such an attack would still be low compared to that in the Erebos Scenario, however, since the consequence would be much lower. For example, in the 2003 Northeast blackout, power was restored to most consumers in less than forty-eight hours and to all users in less than a week. The total economic impact of that attack has been estimated at \$7 billion¹⁵ (in 2003 dollars), about 10 percent of the lowest consequence estimated for Erebos.

¹⁵ *The Economic Impacts of the August 2003 Blackout* (Washington, DC: ELCON: Electricity Consumers Resource Council, February 9, 2004).

Analysis Limitations

Probabilistic risk assessments are subject to a number of known pitfalls. One source of error is the possibility that events treated as independent really have dependencies. The PCR used in this analysis represents a combination of attacker and defender skill. The malware effectiveness represents the robustness of the adversaries' methods (malware, etc.) in the face of random, uncontrollable, detrimental circumstances that arise so shortly before the launch of the grid attack (or outside the attackers' awareness) so that no adjustment is possible. It is not unreasonable to suppose that more skilled adversaries, as evidenced by higher PCRs, might also be able to design more robust malware, although it is not necessarily related. In our analysis, we assume that the skill set leading to higher PCRs is different from the skill set needed to produce robust malware. If high PCR and high malware effectiveness are correlated, then the adversaries could be more effective than indicated by our results.

We also assume that the ability to compromise one plant does not affect the adversaries' ability to compromise another. This assumption of independence could be violated in at least two ways. First, there might be an easier way to compromise, say, plant 2 from the inside of plant 1 than it is to compromise it from the outside. This is certainly true between some pairs of control rooms, but not all. Without data on the frequency of this connection, we assume that there is no connection among control rooms that is easier to exploit than the connection from the outside.

Second, to the extent that plant 1 and plant 2 have similar equipment and operating procedures (quite possible if the plants are owned by the same company), then after compromising plant 1, the adversaries will have a head start on compromising plant 2. We assume that, even if equipment is similar, the specifics of each power plant are different enough to preclude a significant degree of dependence. If there is dependence among a significant number

of control rooms, for either of these reasons, then, again, the adversaries could be more effective than indicated by our results.

The coarse granularity of this analysis and the lack of any real data on attacker skill probably overwhelm any inaccuracy introduced through unmodeled dependencies. Nonetheless, the degree of dependence for these two factors could be measured (see the Implications for Quantitative Vulnerability Assessment section) and adjusted in future analyses.

The lack of data on attacker skill and behavior is accounted for in this analysis by parameterizing the results for different attacker skill levels. This still leaves the open question: how do we know what the probability of this attack is, since it varies from (essentially) 0 to 100 percent depending on our choice of attacker skill? The Unknown Data section describes the need to overcome this limitation to any risk assessment, of any type, by deliberate experimentation.

Finally, for the purpose of quantifying the cyber attack scenario described in *Business Blackout*, this analysis was challenged by the lack of specificity and consistency in the scenario description. This does not negate the results in this report; it simply makes them suspect when used in conjunction with the economic consequences reported in *Business Blackout*.

Conclusions

The specific results from these analyses (for example, the exact percentages of goal achievement) are specific to the region studied and relied on a number of simplifying assumptions. Further, they are subject to the limitations noted in the Analysis Limitations section. Within the limitations of this analysis, however, we can draw a few general conclusions.

Implications for Grid Defense

The grid's overall vulnerability to cyber attack is often extrapolated from specific vulnerabilities of

some cyber components or from specific incidents at certain facilities. At best, these anecdotes can point to the vulnerability of some components of the grid. In our analysis, the interaction of our modeled adversaries' behavior and the underlying distribution of generators, generator sizes, and generator types had a significant impact on the results. This demonstrated the importance of considering any attack scenario or defensive strategy in light of the whole grid environment. We conclude that the vulnerability of the grid is a property of the whole grid system and the distribution of vulnerable assets (e.g., generators of a certain size and type) among plants and within plants.

Within our area of study, we concluded that small improvements in defense can have an outsized impact on the ability to create a large-scale cascading blackout. In all our analyses, reducing the PCR from 20 to 10 percent changed the probability of goal achievement from virtually certain to improbable. Reducing the PCR further to 5 percent made the GAR essentially zero in all cases. This implies that there is a defense that is "good enough"; that is, measures that make it more difficult for an adversary to control grid assets can be very effective even if they are not perfect. Although the exact percentages we found (5 to 20 percent) are specific to the grid makeup and assumptions of our study, the principle may hold in general.

Holding the adversaries to a small PCR will require that the control rooms and defenses are varied; that is, that they have no common mode failure. If every control room in NPCC and RFC uses the same equipment set up in the same way, a single successful attack method can make the adversaries' success rate close to 100 percent, regardless of how good defenses are against other attacks. At high PCRs, the adversaries were invariably successful (and we believe this *is* a result that will hold across any grid segment). We conclude (along with many others) that retaining a significant diversity in hardware, software, and operating procedures in control rooms is an essential element of grid defense.

For the grid segment studied in this report, we conclude that the defense does not have to be applied equally to reduce the potential for a cascading failure. Making a relative handful of large plants and generators very secure can significantly increase the adversaries' effort required to cause a major blackout. In our analysis, when deploying a very effective defense in twenty control rooms out of our target set of 244 (8 percent), the adversaries with lower PCRs were unable to create an effect. Even adversaries with the high PCRs had to compromise fifteen control rooms rather than ten to achieve their goal (50 percent increase).

This result depends strongly on the underlying distribution of generation in the grid. In the region of our analysis, the largest plants had significantly more capacity than the smaller ones; a more equal distribution would make this measure less effective. Further, generation capacity was essential to the adversaries' goal in our study; for some other goal, highly protecting some other set of plants might be more effective. By analyzing a specific region of interest for overall vulnerability to a particular type of attack effect (e.g., cascading failure in the *Business Blackout* case), the appropriate defense can be found.

Although effort (as a function of number of plants compromised) is a consideration for adversaries, highly motivated, well-resourced adversaries would likely make the extra effort required to overcome their inability to compromise the largest plants; however, increased adversary effort is not the only advantage that accrues to the defender. The more control rooms and generators the adversary has to compromise, the more "noise" the campaign will generate, increasing the odds that adversaries will be detected prior to acquiring all the generation needed for a cascading blackout. Cyber situational awareness was not considered in this analysis; however, this analysis shows that effective monitoring and information sharing will work synergistically with other grid cyber defenses that force the adversary to work harder to succeed.

Implications for Quantitative Vulnerability Assessment

Although relatively crude, our analysis yielded interesting quantitative results. For example, we showed that by randomly attacking plants, even with 100 percent success, adversaries would obtain a set of fifty generators with 18,000 megawatts capacity only about 7 percent of the time—improbable. On the other hand, if the set of generators had been slightly larger (sixty-two on average), the same adversaries would control 18,000 megawatts of capacity about half the time. While no adversaries will be 100 percent successful (we hope), and our results are not definitive, we showed that the difference between improbable and inevitable is perhaps not so large as we would like. We were able to show the difference that strategy makes—both for the defender and the adversary. We conclude that valuable insights can be gained by playing off the uncertainties that face any attacker (that is, the random aspects of cyber attack) against the known properties of the grid, even without specific attack vectors or defenses.

The phrase “garbage in, garbage out” applies to any analysis and is especially pertinent to cyber vulnerability assessment where data are sparse. The authors of *Business Blackout* understood the need for better data, saying, “Key requirements will be to enhance the quality of data available and to continue the development of probabilistic modeling for cyber risk.”¹⁶ The remainder of this section examines the data used in our analysis and where and how it could be improved.

Uncertain Data

Cascade-initiating capacity: The most critical uncertain grid attribute used in this analysis is the criterion for starting a cascading failure. This analysis used the single Erebos Scenario criterion of 18,000 megawatts of capacity lost simultaneously.

This was very clearly a “swag” based on a single measurement of hourly peak load in an area that was ambiguously defined. Our analysis did not explore sensitivity to this parameter. Other parameters we did vary had nonlinear effects; we might see the same nonlinearity with cascade capacity. Beyond the absolute value of the needed capacity, we suspect the distribution of lost capacity would also affect the impact on the grid; that is, *where* generation was lost might have a significant impact on the start of a cascade. The *Business Blackout* analysis studied the economic effect of the geographic distribution of lost generation but did not consider it as part of its cyber attack scenario.

The tools needed to gain a better understanding of cascading behavior are available. The electric industry uses models of electricity flow that range in scope from single utilities to an entire interconnection. These models could be used in a systematic study that better defines the possible goal of the adversaries in subsequent cyber studies. A reasonable potential exists that “families” of cascades will emerge that center on certain patterns of outages. If so, then a far more accurate cyber risk assessment, and a better potential for effective defense, would result.

Generator susceptibility: A second important parameter is the number and distribution of vulnerable generators. Because generator damage was an important part of the *Business Blackout* economic analysis, the adversaries in the Erebos Scenario used an Aurora-like attack. Although the preponderance of generators in NPCC and RFC can be simply described as turbines directly connected to the grid, thus inherently susceptible to out-of-sync connection to the grid, a large variety of specific technologies are used to implement the turbine. Some of these may be harder to damage with an Aurora attack than others, due to mechanical coupling or some other factor that was not known by these analysts. Because of the wide range of capacities associated with various types of generators, restricting the adversaries to some subset of them could have a large impact on their goal

¹⁶ Tom Bolt, “Foreword,” *Business Blackout*.

achievement. Further, the Aurora mechanism is not the only way to damage a generator. Considering which generators are susceptible to other damaging attacks would lead to more complete understanding of the potential for an Erebus Scenario to succeed.

Generator protection: The PCR in this analysis included the probability of overcoming protective devices. As shown in Figure 3, a decrease in any of the three factors can reduce the PCR to a safer level. Although there is a “fix” for the specific approach used in the Aurora experiment, we assumed that the adversaries were able to overcome the protective devices in every case; this overstates the risk. A better approach is to understand, at some level, how well grid assets (in this case, generators) are protected from exploitation by cyber means. For example, intelligent electronic devices (IEDs—a generic term for all types of grid control devices with embedded processors) that can be updated over a network are in an entirely different protection class from IEDs that can be updated only through physical access.

In addition to a general understanding of the hardness of IEDs, the diversity of protective devices will have a major impact. If the majority of utilities use the same device and that device has a single exploitable flaw, then the PCR probability factor for overcoming protective devices goes up. Supply chain attacks are another potential means of exploiting a widely used protective device and raising the PCR. This factor could be better estimated if the distribution of IEDs by manufacturer were known. An analysis like this could be further refined by studying the frequency of update and patch downloads.

Grid makeup: Finally, the EIA data used in this analysis were invaluable, but their organization made them difficult to work with, and their inconsistency made their accuracy suspect. The grid owners and operators supply the information by filling out various forms, and the information is then transcribed into the spreadsheets available on the EIA website. Any such process is subject to error. For the purposes of this analysis, the error introduced

through inaccurate EIA data is probably swamped by errors in assumptions. For future analysis, accurate information on generators, types, plants, etc. may become important.

Unknown Data

The application of probabilistic risk assessment to assessing cyber vulnerability has its detractors, largely because the distributions involved are unknown. The ability to create the needed distributions is hampered by the newness of cyber attack as an important threat and the lack of data in a useful form.

Probability of plant compromise: In our analysis, PCR included generator vulnerability and the vulnerability of protective devices. Data exist that could be used to estimate these factors (see the Uncertain Data section). There are no data available to estimate the probability that a control room can be accessed, penetrated, or implanted with malware. In *Business Blackout*, the data on cyber attacks on industrial control systems (ICSs) consist of a list of fourteen anecdotal reports of significant attacks from 1999 to 2015. The report states that “Sharing of cyber attack data and pooling of claims information is a complex issue, but the systemic, intangible, dynamic nature of cyber risk means that all parties involved in managing the risk have an interest in sharing anonymised data on the frequency and severity of attacks.”

Data sharing has important protective value, but it may not have predictive value even over the long term. Attacks on grid systems are relatively rare; more importantly, the range of systems, types of attacks, and attack goals differ so significantly across incidents that they shed little light on future incidents. There has been only one widespread destructive attack involving ICSs, the sum of all our fears, Stuxnet. Data from Stuxnet should be available and would provide important insight into statistical properties of autonomous attacks on air-gapped systems but would shed little light on other modalities. Unless attacks on grid ICS systems become much more common and consistent (an

eventuality not to be desired), amassing the kind of data that insurance companies like Lloyd's of London normally use to set rates and project losses may never be possible. Likewise, government and industry will not be able to understand the cost/benefit ratio of specific enhancements to grid security.

Today, we only know that facilities must be NERC compliant. NERC compliance paints with a very broad brush. For the purposes of a probabilistic study of grid vulnerability, we need much better understanding of the underlying distribution of cyber vulnerability. These are data we must generate for ourselves through experience and experimentation. Today, red teams are often used, with the full cognizance of network owners, to point out specific individual vulnerabilities in networks, which are then remedied. Once the deficiencies are remedied, the red team's work is thought to have no further value. However, used purposefully, red team results can have value beyond any particular network. In work done to develop metrics for the Office of the Secretary of Defense Network and Information Integration's Global Information Grid Information Assurance Portfolio,¹⁷ red teams attacked the same system before and after certain defenses were put in place, to gauge a change in the level of effort involved. In a step beyond this, one or more red teams could be employed continuously to "roam" the grid, not to point out security flaws (although this can certainly be a useful by-product) but rather to assess the overall PCR across the grid as a function of time and various security architectures.

There are two critical differences between this suggested use of a red team and current practice. First, the red teams in the proposed employment are not assessing the vulnerability of any one organization; they are assessing the vulnerability of the grid as a whole. The purpose of their activity is to provide data for compromise statistics.

The second key point is that their assessment is a function of time, both time to compromise a single facility and the period of time that their assessment is valid. Most cybersecurity experts would agree that, in general, determined adversaries will find a way to achieve their goal; however, it may take multiple tries or long-lead-time approaches. Time is on the side of the defenders in that, as time stretches out for the adversaries, they will be subject to changes in the target system that may render previous work useless. There may be a pace at which the adversaries must attack to ultimately achieve their goal. Red teams can compare their success rates to this minimum effective attack rate.

This same variability in the target system, coupled with the continuous enhancement of the attack tool set, leads to the need for a certain repetition rate for red team assessments to remain valid. Measurement of the factors involved in the time element of assessment—for example, the rate of vulnerability discovery, the time to patch release, and the frequency of updates and replacements—can be collected, if the effort is made. These factors will drive the number of red teams that are needed to cover a grid of a certain size.

Probability of malware effectiveness: The second important probability distribution used in this analysis—the probability that, for any particular day and circumstance, the adversaries' malware would have the desired effect—could be more easily and precisely measured using cloud-based network emulations or cyber ranges. Currently, cyber ranges are used in place of real systems to practice attack and defense, basically replicating the red team experience in an artificial environment. Another kind of cyber range is a virtual environment that contains large numbers of replicas of specific systems (for example, a particular IED or human-machine interface system). Malware released in this virtual environment will succeed and fail according to the sensitivity of the malware and the infinitely variable state of any given

¹⁷ Marc Austin and Ryan Layer, "Protect Data and Networks Metrics Framework" (JHU/APL internal memo, February 6, 2008).

system. Probabilistic models of attack success could leverage these data to give more accurate results.

In any case, the data we need to understand our risk and make informed decisions on grid defenses cannot be left to the adversary to supply. Some scheme involving a deliberate data collection effort on the part of the United States is needed.

Implications for Policy

The conclusions derived from this analysis have significant impact on existing policies and grid protection points of view.

Defending the Grid, Not the Grid Components

A major conclusion of this analysis is that vulnerability to grid-wide attacks, such as inducing a cascading blackout, is a property of the entire grid, not of any single facility or asset. There are potential standards that would improve overall grid resilience while having little or no benefit to any individual asset owner. For example, one way to achieve an overall grid PCR of 10 percent is to require every plant to have a 10 percent PCR, a stringent requirement that may be costly and difficult to achieve. Another way to achieve a 10 percent PCR is to ensure enough diversity across the grid that, at any given time, adversaries can compromise only 10 percent of the plants using a small set of exploits. This might be achievable at virtually no cost, since the market drives equipment costs across vendors to equalize. There currently is no mechanism for a government or industry agency to set standards for a grid segment, rather than for every asset, although the former might be more cost effective.

Our analysis showed that the defense of some plants is more important to the overall security of the grid than the defense of others is, another example of grid-wide defense. If asked to radically improve the security of these key plants, their owners might ask why they should bear the cost of a defense that would benefit everyone. If the conclusions from future

analyses of this type were operationalized, a debate on cost sharing for this element of grid defense will ensue, involving the electric industry, government at federal, state, and local levels; and commercial and residential consumers. It seems reasonable that something like the economic analysis performed in *Business Blackout* could provide a starting point for determining *cui bono*.

Government–Industry Relationships

The use of red teams as suggested in the Unknown Data section has broad implications. We see no precedent for such an activity. On the other hand, the circumstances of cyber attack—carried out remotely and possibly without attribution—are unique. The introduction of cyber operations to warfare, in particular the potential for attack on civilian populations, may necessitate rethinking conventional approaches.

First, in the course of its work, the red team itself may cause a power outage. With proper precautions, the possibility is small, but it cannot be eliminated. Such an outage imposes a cost to both the electric industry and the electricity user. The cost to the industry comes both in the form of lost revenue and potential fines for failing to meet government-imposed standards for reliability. In a chicken-and-egg-like situation, this cost may be trivial compared to the risk of losing far more in a significant adversarial cyber attack; however, the red team activity is needed to assess the risk of a significant adversarial cyber attack for comparison.

Other considerations may outweigh the potential disruption of service caused by a red team. Although the concept of public–private partnerships is evolving, government agencies like the Federal Energy Regulatory Commission and standards-setting groups like NERC have a regulatory and enforcement role. A policy for employing red teams in a discovery mode would have to deal with the very fine line between beneficial regulatory enforcement and punitive regulatory enforcement, especially

for quantities like PCR, which apply to the grid as a whole and which no single company can improve by itself. Although the government is clearly better positioned to play the part of a nation-state adversary, the industry might want to employ its own red teams, to keep the individual results private. Aggregate grid-wide results could be released to the government to promote public understanding of the risk. Even if red teams were in the employ of the electric industry, however, the data they collect could be the basis for lawsuits if individual company deficiencies (as opposed to grid-wide deficiencies) are not remedied. Public policy would be needed to protect industry from private-citizen reprisals for real or imagined injury.

Data Access

We conclude that better data are required for definitive quantitative vulnerability analysis that could form the basis for policy. Some of the data are (in principle) known and could be collected from the

electric industry, much as the data that appear on the EIA website are today, at some recurring cost. Which entity should bear this cost is another policy question.

Some data required for research to improve grid security already exist, held both by industry and the government; however, researchers are often denied access. The reluctance to share these data is based on a reasonable fear that the data would be helpful to adversaries planning to attack the grid. On the other hand, breaches at the Office of Personnel Management on the government side, and Sony and Google on the industry side, have demonstrated that adversaries can almost certainly get access to these data if they want to. Only the researcher attempting to improve the defense of the United States cannot. Presumably, with decades of experience in protecting classified data, we could establish some uniform set of standards for grid data that would benefit defense without giving adversaries better access to the data than they already have.

Appendix A Makeup of the Grid Segment Used in the Analysis

The US Energy Information Administration (EIA) data set associates 2,426 power plants with about 15,745 generators in the Northeast Power Coordinating Council (NPCC) and ReliabilityFirst Corporation (RFC) regions. Because the underlying distribution of plant and generator size contributed to the analysis result, this appendix breaks down the regional grid in various ways.

The distribution of generators among plants is uneven. Figure A-1 shows the distribution of generators across the plants in the region. Note that the figure is plotted on a logarithmic scale, since plants with one generator dominate plants with more.

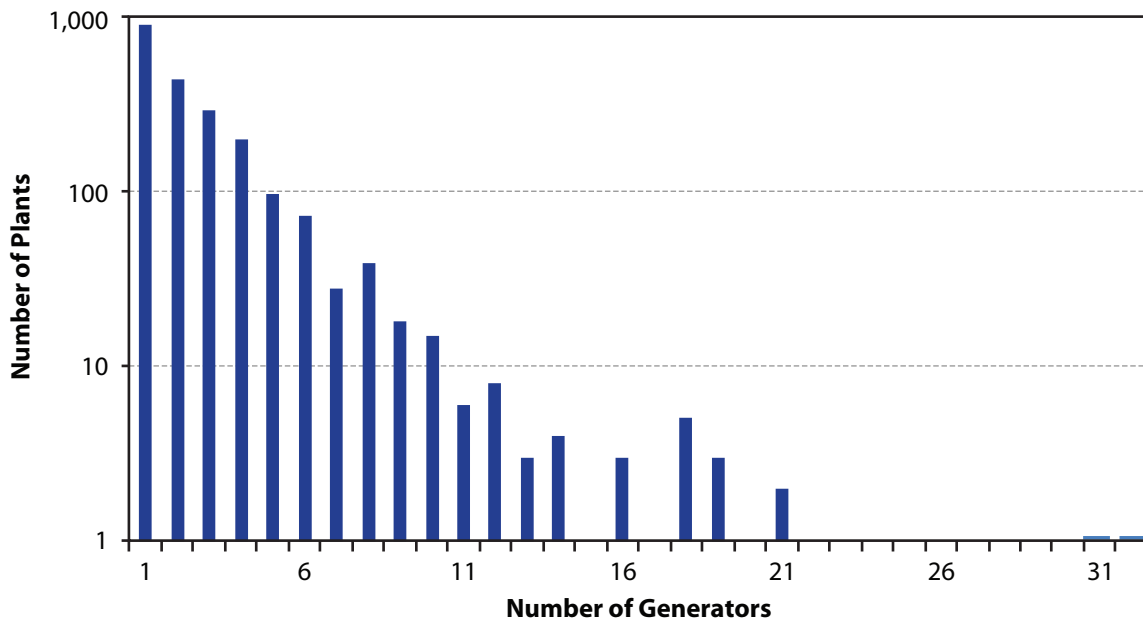


Figure A-1. Number of Generators per Plant in NPCC and RFC (Excluding 48 in Mountain View and 73 in Edison Sault)

Capacity varies widely among generators. Figure A-2 shows the frequency of generators of a given capacity across the region. This figure is also plotted on a logarithmic scale, with an expanded number of capacity bins below 100 megawatts, because small generators are far more prevalent than large generators.

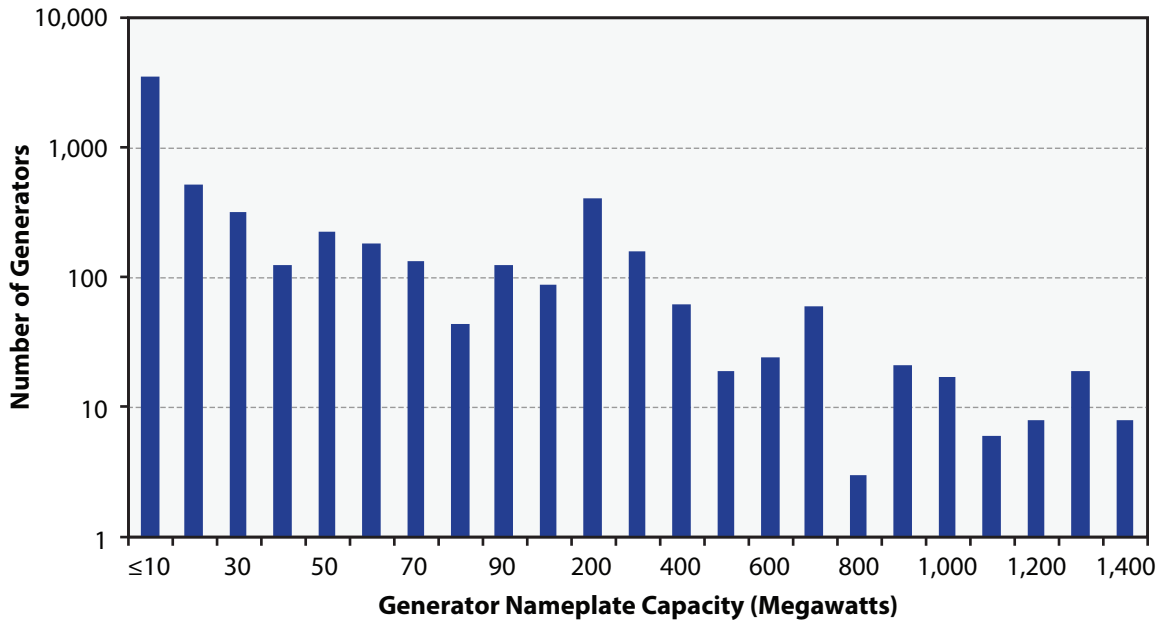


Figure A-2. Number of Generators by Capacity in NPCC and RFC

There are at least twenty-three different combinations of generator technology and fuel types ranging from fission to land-fill gases. Figure A-3 shows the number of generators by type in the NPCC and RFC regions.

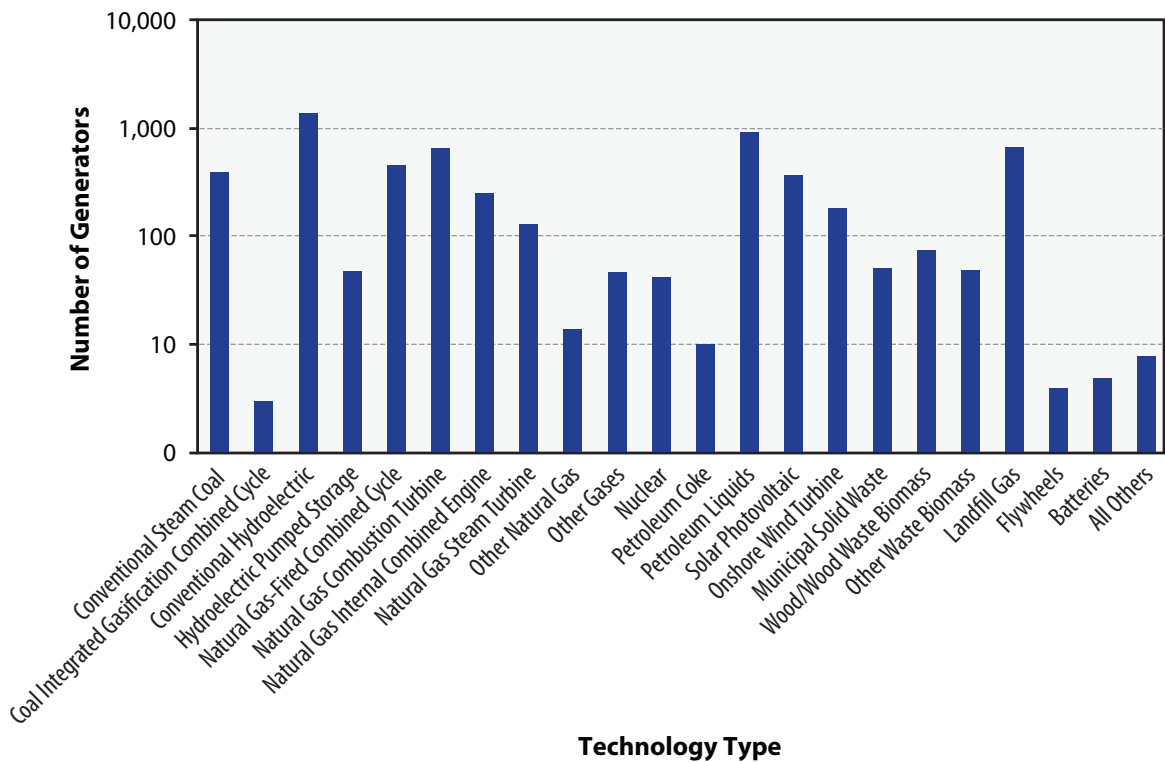


Figure A-3. Number of Generators by Type in NPCC and RFC

The capacity of a generator of a given type could also vary widely. Figure A-4 shows the range of generator capacity for each type in NPCC and RFC.




Figure A-4. Generator Capacity by Type

For the target set used in most of the analysis (plants with generators with at least 100 megawatts, excluding nuclear, solar, and wind generators), the total capacity of individual plants is shown in descending order in Figure A-5.

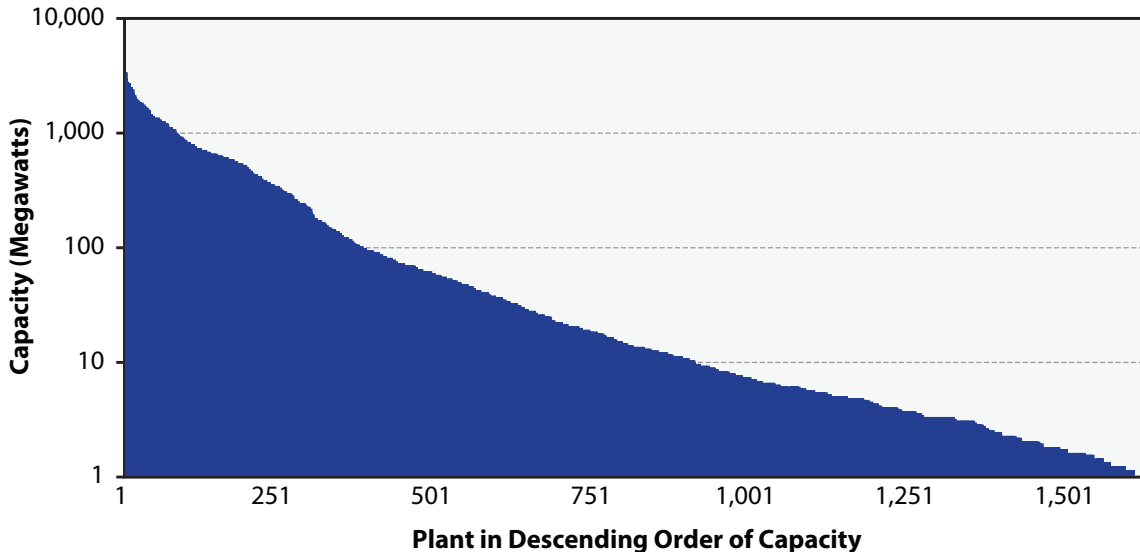


Figure A-5. Total Plant Capacities in Descending Order (Nuclear, Solar, and Wind Excluded)

Total plant capacity is only weakly correlated with the number of generators in the plant, if at all. Figure A-6 shows the number of generators in each plant in our target set plotted against the plant’s total capacity.




Figure A-6. Plant Capacity vs. Number of Generators in the Plant

Plants contained widely varying sizes of generators, as well as number of generators. Figure A-7 shows the six largest plants in our target set with total plant capacity, plotted with the sizes of the generators they contain. This idiosyncrasy greatly depressed the adversaries’ Goal Achievement Rate (GAR) for the Erebus Scenario if we assumed that the adversaries damaged all the generators within each plant. For example, if they were lucky enough to compromise the first, third, and fourth plants shown in Figure A-7, they could damage eleven generators and take about 9,000 megawatts of capacity out of the grid. On the other hand, if they were unlucky enough to compromise the second, fifth, and sixth plants in the figure, they could damage forty-two generators, but still only take about 9,000 megawatts out of the grid. In each case, after compromising three plants, they control about half the generation they need; however, in the second case, they are much more likely to hit the fifty-generator limit before getting to 18,000 megawatts.

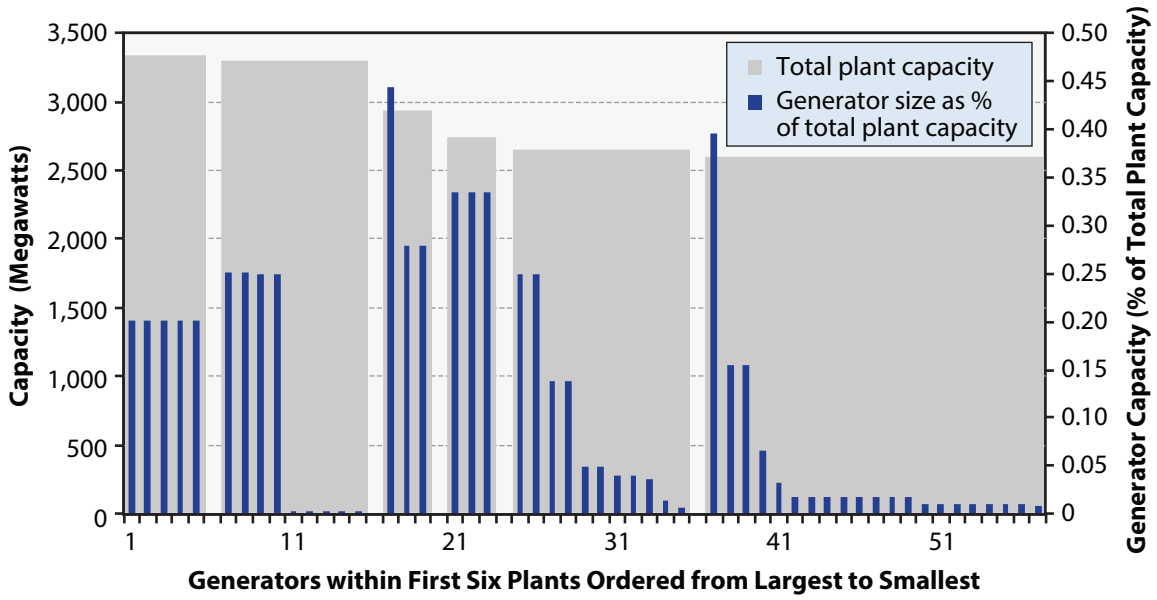


Figure A-7. Unequal Distribution of Generator Capacity within Plants

Appendix B List of Plants Used in the Analysis

This list of 244 plants with generators with at least 100 megawatts formed the basis for the targets sets we used in our analysis. The base target set has a generation capacity of about 202,000 megawatts, out of a total of about 270,000 megawatts in the US portions of the NPCC and RFC regions.

Table B-1. Plants with Generators with at Least 100 Megawatts of Capacity

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
384	Joliet 29	RFC	PJM	1320.0
546	Montville Station	NPCC	ISNE	414.9
547	Northfield Mountain	NPCC	ISNE	940.0
562	Middletown	NPCC	ISNE	767.9
568	Bridgeport Station	NPCC	ISNE	400.0
593	Edge Moor	RFC	PJM	622.8
594	Indian River Generating Station	RFC	PJM	445.5
599	McKee Run	RFC	PJM	113.6
602	Brandon Shores	RFC	PJM	1370.0
874	Joliet 9	RFC	PJM	360.4
876	Kincaid Generation LLC	RFC	PJM	1319.0
879	Powerton	RFC	PJM	1785.6
883	Waukegan	RFC	PJM	681.7
884	Will County	RFC	PJM	598.4
988	Tanners Creek	RFC	PJM	1100.1
990	Harding Street	RFC	MISO	824.0
991	Eagle Valley (IN)	RFC	MISO	113.6
994	AES Petersburg	RFC	MISO	2146.7
995	Bailly	RFC	MISO	603.5
997	Michigan City	RFC	MISO	540.0
1001	Cayuga	RFC	MISO	1183.0
1004	Edwardsport	RFC	MISO	804.5
1008	R Gallagher	RFC	MISO	300.0
1010	Wabash River	RFC	MISO	860.2

(continued)

Table B-1 (continued)

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
1012	F B Culley	RFC	MISO	368.9
1043	Frank E Ratts	RFC	MISO	233.2
1353	Big Sandy	RFC	PJM	1096.8
1507	William F Wyman	NPCC	ISNE	746.0
1552	C P Crane	RFC	PJM	399.8
1553	Gould Street	RFC	PJM	103.5
1554	Herbert A Wagner	RFC	PJM	1042.5
1556	Perryman	RFC	PJM	192.0
1560	Westport	RFC	PJM	121.5
1564	Vienna Operations	RFC	PJM	162.0
1571	Chalk Point LLC	RFC	PJM	2502.0
1572	Dickerson	RFC	PJM	914.0
1573	Morgantown Generating Plant	RFC	PJM	1252.0
1588	Mystic Generating Station	NPCC	ISNE	2361.4
1595	Kendall Square Station	NPCC	ISNE	186.2
1599	Canal	NPCC	ISNE	1165.0
1619	Brayton Point	NPCC	ISNE	1600.1
1642	NAEA Energy Massachusetts LLC	NPCC	ISNE	113.6
1695	B C Cobb	RFC	MISO	312.6
1702	Dan E Karn	RFC	MISO	1946.3
1710	J H Campbell	RFC	MISO	1560.8
1713	Ludington	RFC	MISO	1978.8
1720	J C Weadock	RFC	MISO	312.6
1723	J R Whiting	RFC	MISO	345.4
1733	Monroe (MI)	RFC	MISO	3279.6
1740	River Rouge	RFC	MISO	933.2
1743	St Clair	RFC	MISO	1547.0
1745	Trenton Channel	RFC	MISO	775.5
1832	Erickson Station	RFC	MISO	154.7

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts)^a
2364	Merrimack	NPCC	ISNE	459.2
2378	B L England	RFC	PJM	339.6
2393	Gilbert	RFC	PJM	296.0
2398	Bergen Generating Station	RFC	PJM	1400.8
2403	PSEG Hudson Generating Station	RFC	PJM	659.7
2406	PSEG Linden Generating Station	RFC	PJM	1355.6
2408	PSEG Mercer Generating Station	RFC	PJM	768.0
2411	PSEG Sewaren Generating Station	RFC	PJM	567.4
2480	Danskammer Generating Station	NPCC	NYIS	386.5
2490	Arthur Kill Generating Station	NPCC	NYIS	877.5
2493	East River	NPCC	NYIS	716.2
2500	Ravenswood	NPCC	NYIS	1997.0
2511	E F Barrett	NPCC	NYIS	376.0
2516	Northport	NPCC	NYIS	1548.0
2517	Port Jefferson	NPCC	NYIS	376.0
2527	Greenidge Generation LLC	NPCC	NYIS	112.5
2535	Cayuga Operating Company	NPCC	NYIS	322.5
2539	Bethlehem Energy Center	NPCC	NYIS	893.1
2549	C R Huntley Generating Station	NPCC	NYIS	400.0
2554	Dunkirk Generating Plant	NPCC	NYIS	435.2
2594	Oswego Harbor Power	NPCC	NYIS	1803.6
2625	Bowline Point	NPCC	NYIS	1242.0
2691	Blenheim Gilboa	NPCC	NYIS	1000.0
2693	Robert Moses Niagara	NPCC	NYIS	2429.1
2828	Cardinal	RFC	PJM	1880.4
2831	Dicks Creek	RFC	PJM	100.0
2832	Miami Fort	RFC	PJM	1278.0
2835	FirstEnergy Ashtabula	RFC	PJM	256.0
2836	Avon Lake	RFC	PJM	680.0

(continued)

Table B-1 (continued)

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
2837	FirstEnergy Eastlake	RFC	PJM	369.0
2838	FirstEnergy Lake Shore	RFC	PJM	256.0
2840	Conesville	RFC	PJM	1729.3
2843	Picway	RFC	PJM	106.2
2847	Frank M Tait	RFC	PJM	209.6
2850	J M Stuart	RFC	PJM	2440.8
2866	FirstEnergy W H Sammis	RFC	PJM	2455.6
2872	Muskingum River	RFC	PJM	1529.4
2878	FirstEnergy Bay Shore	RFC	PJM	150.5
2880	Richland	RFC	PJM	405.0
3096	Brunot Island	RFC	PJM	144.0
3113	Portland (PA)	RFC	PJM	583.0
3118	Conemaugh	RFC	PJM	1872.0
3122	Homer City Generating Station	RFC	PJM	2012.0
3130	Seward (PA)	RFC	PJM	585.0
3131	Shawville	RFC	PJM	626.0
3136	Keystone	RFC	PJM	1872.0
3138	New Castle Plant	RFC	PJM	250.0
3140	Brunner Island	RFC	PJM	1616.1
3148	PPL Martins Creek	RFC	PJM	1701.0
3149	PPL Montour	RFC	PJM	1757.9
3161	Eddystone Generating Station	RFC	PJM	782.0
3164	Muddy Run	RFC	PJM	1072.0
3236	Manchester Street	NPCC	ISNE	375.0
3775	Clinch River	RFC	PJM	712.5
3776	Glen Lyn	RFC	PJM	337.5
3780	Smith Mountain	RFC	PJM	415.5
3935	John E Amos	RFC	PJM	2932.6
3936	Kanawha River	RFC	PJM	439.2

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts)^a
3938	Philip Sporn	RFC	PJM	610.0
3943	FirstEnergy Fort Martin Power Station	RFC	PJM	1152.0
3944	FirstEnergy Harrison Power Station	RFC	PJM	2052.0
3947	Kammer	RFC	PJM	712.5
3948	Mitchell (WV)	RFC	PJM	1632.6
3954	Mt Storm	RFC	PJM	1662.4
4040	Port Washington Generating Station	RFC	MISO	1208.8
4041	South Oak Creek	RFC	MISO	1240.0
4042	Valley (WI)	RFC	MISO	272.0
5083	Cumberland (NJ)	RFC	PJM	131.8
6004	FirstEnergy Pleasants Power Station	RFC	PJM	1368.0
6018	East Bend	RFC	PJM	669.3
6019	W H Zimmer	RFC	PJM	1425.6
6031	Killen Station	RFC	PJM	660.6
6034	Belle River	RFC	MISO	1395.0
6035	Greenwood (MI)	RFC	MISO	815.4
6081	Stony Brook	NPCC	ISNE	105.0
6082	Somerset Operating Co LLC	NPCC	NYIS	655.1
6085	R M Schahfer	RFC	MISO	2201.4
6094	FirstEnergy Bruce Mansfield	RFC	PJM	2741.1
6113	Gibson	RFC	MISO	3339.5
6137	A B Brown	RFC	MISO	530.4
6156	New Haven Harbor	NPCC	ISNE	460.0
6166	Rockport	RFC	PJM	2600.0
6170	Pleasant Prairie	RFC	MISO	1233.2
6213	Merom	RFC	MISO	1080.0
6264	Mountaineer	RFC	PJM	1300.0
6522	Yards Creek	RFC	PJM	453.0
6705	Warrick	RFC	MISO	822.8

(continued)

Table B-1 (continued)

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
7153	Hay Road	RFC	PJM	1193.0
7288	Sherman Avenue	RFC	PJM	112.8
7314	Richard M Flynn	NPCC	NYIS	108.0
7835	NAEA Rock Springs LLC	RFC	PJM	772.6
7872	Robert P Mone Plant	RFC	PJM	594.0
8002	Newington	NPCC	ISNE	414.0
8005	Bear Swamp	NPCC	ISNE	600.0
8006	Roseton Generating Facility	NPCC	NYIS	1242.0
8102	General James M Gavin	RFC	PJM	2600.0
8225	FirstEnergy Seneca	RFC	PJM	440.0
8226	Cheswick Power Plant	RFC	PJM	637.0
8906	Astoria Generating Station	NPCC	NYIS	1330.0
10043	Logan Generating Company LP	RFC	PJM	242.3
10143	Colver Power Project	RFC	PJM	118.0
10307	Bellingham Cogeneration Facility	NPCC	ISNE	386.1
10308	Sayreville Cogeneration Facility	RFC	PJM	430.2
10495	Rumford Cogeneration	NPCC	ISNE	102.6
10566	Chambers Cogeneration LP	RFC	PJM	285.0
10676	AES Beaver Valley Partners Beaver Valley	RFC	PJM	114.0
10678	AES Warrior Run Cogeneration Facility	RFC	PJM	229.0
10725	Selkirk Cogen	NPCC	NYIS	148.4
10745	Midland Cogeneration Venture	RFC	MISO	790.0
50006	Linden Cogen Plant	RFC	NYIS	212.5
50243	Bucksport Generation LLC	NPCC	ISNE	186.8
50733	Gary Works	RFC	MISO	161.0
50888	Northampton Generating Company LP	RFC	PJM	134.1
54547	Sithe Independence Station	NPCC	NYIS	1086.1
54785	Grays Ferry Cogeneration	RFC	PJM	135.0
54805	Milford Power LP	NPCC	ISNE	128.9

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts)^a
54914	Brooklyn Navy Yard Cogeneration	NPCC	NYIS	242.0
55011	LSP-Whitewater LP	RFC	MISO	283.5
55026	Dighton Power Plant	NPCC	ISNE	200.0
55041	Berkshire Power	NPCC	ISNE	289.0
55042	Bridgeport Energy Project	NPCC	ISNE	520.0
55048	Tiverton Power Plant	NPCC	ISNE	179.3
55068	Maine Independence Station	NPCC	ISNE	550.2
55079	Millennium Power	NPCC	ISNE	360.0
55087	Zeeland Generating Station	RFC	MISO	968.2
55088	Dearborn Industrial Generation	RFC	MISO	760.0
55100	Rumford Power, Inc	NPCC	ISNE	179.4
55107	Energy Rhode Island State Energy LP	NPCC	ISNE	596.0
55109	Rocky Road Power LLC	RFC	PJM	374.0
55126	Milford Power Project	NPCC	ISNE	578.0
55131	Kendall County Generation Facility	RFC	PJM	1256.0
55135	Alliant Energy Neenah	RFC	MISO	371.0
55149	Lake Road Generating Plant	NPCC	ISNE	840.0
55170	Granite Ridge	NPCC	ISNE	790.0
55188	Cordova Energy	RFC	PJM	611.2
55193	Ontelaunee Energy Center	RFC	PJM	728.0
55198	Riverside Generating LLC	RFC	PJM	1150.0
55199	Elwood Energy LLC	RFC	PJM	1728.0
55211	ANP Bellingham Energy Project	NPCC	ISNE	578.0
55212	ANP Blackstone Energy Project	NPCC	ISNE	578.0
55224	Wheatland Generating Facility	RFC	MISO	540.0
55231	Liberty Electric Power Plant	RFC	PJM	614.0
55236	Lee Energy Facility	RFC	PJM	814.4
55238	NRG Rockford I	RFC	PJM	316.0
55239	Red Oak Power LLC	RFC	PJM	821.1

(continued)

Table B-1 (continued)

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
55259	Whiting Clean Energy	RFC	MISO	576.8
55270	Jackson Power Facility	RFC	MISO	210.0
55279	Aurora	RFC	PJM	849.0
55294	Westbrook Energy Center Power Plant	NPCC	ISNE	563.9
55296	Calumet Energy Team LLC	RFC	PJM	312.8
55297	New Covert Generating Facility	RFC	MISO	1176.0
55298	Fairless Energy Center	RFC	PJM	1338.0
55317	Fore River Generating Station	NPCC	ISNE	872.2
55337	PPL Ironwood LLC	RFC	PJM	777.6
55347	Armstrong	RFC	PJM	688.0
55348	Troy Energy LLC	RFC	PJM	688.0
55349	Pleasants Energy LLC	RFC	PJM	344.0
55350	Dresden Energy Facility	RFC	PJM	678.3
55364	Sugar Creek Power	RFC	MISO	619.4
55375	Astoria Energy	NPCC	NYIS	592.0
55392	Zion Energy Center	RFC	PJM	596.7
55397	Washington Energy Facility	RFC	PJM	714.9
55401	Rolling Hills Generating	RFC	PJM	977.5
55402	Renaissance Power LLC	RFC	MISO	680.0
55405	Athens Generating Plant	NPCC	NYIS	1323.0
55438	Elgin Energy Center LLC	RFC	PJM	540.0
55502	Lawrenceburg Energy Facility	RFC	PJM	1232.0
55503	AEP Waterford Facility	RFC	PJM	921.6
55516	Fayette Energy Facility	RFC	PJM	644.1
55524	York Energy Center	RFC	PJM	560.0
55661	EP Newington Energy LLC	NPCC	ISNE	605.5
55667	Lower Mount Bethel Energy	RFC	PJM	651.6
55690	Bethlehem Power Plant	RFC	PJM	1300.0
55701	Fremont Energy Center	RFC	PJM	739.5

NERC Plant ID	Plant Name	NERC Region	Balancing Area (BA)	Total Capacity (Megawatts) ^a
55710	FirstEnergy Allegheny Energy Units 3 4 & 5	RFC	PJM	556.0
55736	Hanging Rock Energy Facility	RFC	PJM	1288.2
55801	FPL Energy Marcus Hook LP	RFC	PJM	836.1
55936	NRG Rockford II Energy Center	RFC	PJM	168.0
55938	NAEA Ocean Peaking Power LLC	RFC	PJM	383.0
55976	Hunterstown Power Plant	RFC	PJM	898.0
56031	Fox Energy Center	RFC	MISO	618.8
56068	Elm Road Generating Station	RFC	MISO	1402.6
56196	500 megawatts CC	NPCC	NYIS	528.0
56234	Caithness Long Island Energy Center	NPCC	NYIS	348.9
56259	Empire Generating Co LLC	NPCC	NYIS	653.7
56671	Longview Power LLC	RFC	PJM	807.5
56798	Kleen Energy Systems Project	NPCC	ISNE	693.0
56808	Virginia City Hybrid Energy Center	RFC	PJM	668.0
56963	West Deptford Energy Station	RFC	PJM	754.0
57664	Astoria Energy II	NPCC	NYIS	650.0
57842	Wabash Valley Power IGCC	RFC	MISO	304.5

ISNE, Independent System Operator New England; MISO, Midcontinent Independent System Operator; NERC, North American Electric Reliability Corporation; NPCC, Northeast Power Coordinating Council; NYIS, New York Independent System Operator; PJM, Pennsylvania-Jersey-Maryland; RFC, ReliabilityFirst Corporation.

^a The total capacity includes only those generators greater than 100 megawatts in each plant.

Acknowledgments

We thank Dr. Paul Stockton, a fellow at the Johns Hopkins University Applied Physics Laboratory, for his governmental insight and Mr. Steve Nauman, a senior industry consultant, for his operational insight. Both contributed to making this a better product. We also thank Dr. David Helmer, Dr. James Scouras, and Ms. Donna Gregg, whose careful review and perceptive comments added to both the correctness and the clarity of the final document. We especially thank Ms. Christine Fox, assistant director for policy and analysis at the Johns Hopkins University Applied Physics Laboratory, and Mr. Matthew Schafer, managing executive of the Laboratory's National Security Analysis Department, for their unflagging support and deep commitment to this work.

About the Authors

Susan Lee is a member of the Principal Professional Staff at the Johns Hopkins University Applied Physics Laboratory. Ms. Lee heads an internal program that performs analysis on issues facing the critical infrastructure of the United States. In part, this program develops innovative analysis techniques to improve understanding of cyber threats to critical infrastructure. Prior to her current role, Ms. Lee spent nearly two decades as chief scientist for the Laboratory sector addressing challenges in cyber operations, special operations, and homeland protection. She holds master's degrees in technical management and computer science from Johns Hopkins University and a bachelor's degree in physics from Duke University.

Michael Moskowitz is a member of the Senior Professional Staff at the Johns Hopkins University Applied Physics Laboratory. Mr. Moskowitz conducts operations analysis on a variety of topics facing the Department of Defense and the Department of Homeland Security. Prior to joining the Laboratory, Mr. Moskowitz spent a decade at the Center for Naval Analyses, including two field rotations supporting the United States Marine Corps. Mr. Moskowitz has a master's degree in applied economics from Johns Hopkins University and a bachelor's degree in economics, with a minor in statistics, from Boston University.

Jane Pinelis is a member of the Senior Professional Staff at the Johns Hopkins University Applied Physics Laboratory. Dr. Pinelis conducts statistical and operations analysis supporting the Department of Defense, the Department of Homeland Security, and Johns Hopkins Hospital. Prior to joining the Laboratory, Dr. Pinelis spent seven years at the Center for Naval Analyses, including field rotations supporting operational testing in the Navy and the Marine Corps, as well as wargaming in the Marine Corps. Dr. Pinelis has a Ph.D. in statistics, a master's degree in statistics, and a bachelor's degree in statistics and economics, with a minor in mathematics, all from the University of Michigan.



JOHNS HOPKINS
APPLIED PHYSICS LABORATORY