SAND2020-4359B

Resilient Design Considerations for Electrical Power and Other Critical Infrastructure Systems

By

Meghan Galiardi and Eric D. Vugrin
Sandia National Laboratories[a]
Albuquerque, NM 87185-7057

For consideration to be included in
Resilient Control Architectures and Power Systems

Abstract

Previous chapters focus on resilient architectures for the electrical power grid. Emphasis on this "uniquely critical" infrastructure system is merited, and other critical infrastructures can also benefit from design of resilient control systems. This chapter discusses resilient design considerations that generally apply across a broad spectrum of critical infrastructures. The chapter introduces four resilient design capacities, that is, fundamental system attributes that contribute to or detract from resilient operations. The chapter also discusses design issues and system constraints that often need to be considered when balancing the capacities in resilient designs.

Learning Objectives

- Introduce and define four resilient design capacities that commonly apply across electrical power and other critical infrastructure systems
- Provide examples of tangible resilience enhancement features that contribute to each of the four capacities
- Describe additional considerations, issues, constraints, and tradeoffs that commonly arise during the resilient infrastructure design process.

Body

Introduction

The previous chapters focus on resilient architectures, evaluation, control, and associated challenges for the electrical power grid. The emphasis on this "uniquely critical infrastructure"[b] is deserved because the strong dependence of all other critical infrastructure and, more generally, modern society on

---

[a] This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

[b] In 2013 Presidential Policy Directive 21 (PPD-21) Critical Infrastructure Security and Resilience designated energy and communications infrastructure as *uniquely critical* "due to the enabling functions they provide across all critical infrastructure sectors".

electrical power. The loss of power for extended and sometimes even brief periods of time can have significant economic, health, and security impacts to communities, regions, and industries.

When considering "how" to design resilient infrastructures, the electrical power grid provides many options and examples for doing so; however, many examples of good (and sometimes bad) resilient design features can be found in other infrastructure systems. With some experience and study across many infrastructure systems, one can start to identify resilient design commonalities that exist, at a conceptual level, across electrical power, and across other critical infrastructure systems. Furthermore, regardless of the infrastructure system under consideration, designers must decide which designs and architectures best achieve resilience objectives while satisfying a myriad of other infrastructure goals and constraints.

This chapter discusses issues that infrastructure owners and operators commonly face when attempting to design resilient infrastructure systems. The discussion herein will apply to both control systems for electrical power systems and other various infrastructure systems. The chapter begins with an introduction to and definition of four resilient design capacities. At an abstract level, these capacities describe fundamental system properties that can determine the resilience of a system. The chapter further provides examples of resilience enhancement features that contribute to one or more capacities. These examples describe tangible technologies, designs, and procedures that can be implemented to contribute to the overall resilience of the system. The chapter concludes with a discussion of common tradeoffs that designers must consider when planning for resilience. Rarely, if ever, is resilience the primary and sole consideration of designers. Rather, designers must find an acceptable balance between resilience and cost, regulatory, environmental, and other priorities.

Please note that this chapter does not provide a detailed, step-by-step process on how to optimally design a resilient system in a constrained decision space; the body of resilient design research is not so mature that doing so is possible yet. Still, by providing a discussion on resilient design capacities and tradeoffs, this chapter will familiarize the reader with many of the issues and challenges that designers commonly face.

Resilient Design Capacities

Vugrin et al. (2011) first defined resilient design capacities to be fundamental system properties that collectively contribute to or detract from a system's overall resilience [1]. At an abstract level, these four capacities are a part of all infrastructure systems. Vugrin et al. further defined resilience enhancement features to be the tangible, infrastructure design features that are implemented and put in place with the intent of improving the resilience of infrastructure systems to a variety of threats.

Vugrin et al. first introduced three resilient design capacities: absorptive, adaptive, and restorative. Since the capacities were first introduced, cyber-physical systems, including the electrical power grid and other infrastructure, have increasingly become the targets of stealthy cyber attacks (e.g., Dragos Inc. [2,3]). These instances have highlighted the need for a fourth capacity that describes a system's ability to detect threats and monitor operations to foster increased resilience, which from hereon is

referred to as the anticipative capacity. This chapter describes and expands upon Vugrin et al.'s initial specification of resilient design capacities.[c]

Anticipative Capacity

A system's anticipative capacity is its ability to identify, categorize, predict, and provide advanced warning of threats to enable a rapid, proactive response. Alone, the anticipative capacity has no direct impact on mitigating threats. However, features contributing to the anticipative capacity can act as a catalyst for other capacities to make absorptive, adaptive, and restorative features more effective. Anticipative features are typically most effective when implemented before the onset of threats and when they can recognize threats before negative consequences are realized.

Examples of resilience enhancement features that contribute to a system's restorative capacity include:

- *Intrusion Detection Systems (IDSs)*: IDSs monitor operations to detect possible threats and alert operators. Commonly used in cyber applications, IDSs may monitor network traffic and other behaviors to detect threats. In physical systems, IDSs may include infrared sensors, vibration detectors, cameras, and other monitors to detect physical threats.
- *State of Health Monitoring*: Whereas IDSs detect and report the presence of a threat, state-of-health monitoring reports when system functions are failing and not operating properly. Such alerting can notify operators to investigate and remediate the cause.
- *Stress Testing*: During stress testing, system operators voluntarily expose their systems to threats (in a controlled manner) so that the operators can identify potential risks and how to effectively address them before the threats occur in an uncontrolled environment. These tests can take on many forms; they can be simple thought exercises or drills. Sometimes, red teams are engaged to probe defenses and find vulnerabilities. Netflix has gone so far as to use the Chaos Monkey tool to randomly terminate instances on operational systems, which ensures high reliability of its streaming platforms [4].
- *Threat Intelligence*: Greater awareness of the threat environment and information sharing can notify system operators of potential threats, giving operators time to prepare for threats and possibility eliminate vulnerabilities.

Absorptive Capacity

The absorptive capacity is the extent to which a system can automatically absorb or withstand the impacts of a threat and minimize negative consequences with relatively low levels of effort. Ideally, system features contributing to the absorptive capacity are installed prior to the realization of a threat and take little effort during and after the threat realization to provide benefit to system operations. Hence, system features contributing to absorptive capacity can sometimes be sufficient for mitigating the effects of a lesser threat.

---

[c] Several resilience frameworks include concepts analogous to these four resilience capacities. For example, PPD-21 mentions *preparing, withstanding, adapting, and recovering* in its definition of resilience. NIST's cyber resilience systems engineering standard *NIST SP 800-160* similarly defines anticipate, withstand, recover, and adapt as resilience goals [5]. The four concepts in these two frameworks, as well as those mentioned in others, map reasonably well to the absorptive, adaptive, restorative, and anticipative capacities and can be used almost interchangeably.

Examples of resilience enhancement features that contribute to a system's absorptive capacity include:

- Compartmentalization (segregation). In the event that a portion of the control systems are compromised, compartmentalization can limit the extent to which adversaries can negatively affect system operations in other portions of the system.
- Decentralization: Spreading system functions and operations to different locations can limit damage from natural disasters. Against human attackers, decentralization can increase resources and time required to cause detrimental effects.
- Excess Capacity: Excess capacity can facilitate continued operations if some portion of the infrastructure system is rendered nonfunctional or is stressed. Excess capacity in electric power generation permits ramping up when demand is high. Excess bandwidth is a common design feature in communication networks to handle spikes in demand or denial-of-service attacks.
- Redundancy: Inclusion of multiple devices and components that can execute the same operation can mitigate consequences if a device fails or is attacked. Triple modular redundancy is a common design aspect in safety and reliability systems.
- Diversity: Redundancy may not be effective against a cyber attack. (If one device contains a vulnerability, another identical device likely includes that same vulnerability.) Diversity in ecosystems can lead to increased sustainability; researchers have postulated that diversity in design and components can potentially enhance resilience.
- Storage: Energy storage is becoming a more-commonly used technology to protect against disruptions of power systems. This approach is also used in manufacturing and supply chains, agriculture, and other infrastructure systems.

Adaptive Capacity

The adaptive capacity is the extent to which a system is capable of adapting and changing to non-standard operating modes in an attempt to overcome the effects of a threat. Activation of adaptive capacity features typically requires an active, dynamic effort and may incur greater costs (measured in terms of money, manpower, time, and other resources) and be less efficient to implement than proceeding according to normal operating procedures. The benefit of their use decreases negative consequences and increases loss avoidance. Because of the greater costs, system operators will generally focus on first activating absorptive capacities features; features contributing to the adaptive capacity will only be activated if the absorptive capacity alone is insufficient to mitigate threat effects or if the perceived consequences of not activating them are deemed to exceed the cost of implementing non-standard operating modes. Adaptive measures are generally implemented temporarily, with systems returning to normal operations when the threat is overcome.

Examples of resilience enhancement features that contribute to a system's adaptive capacity include:

- Rerouting: Using alternative transport pathways is a common approach when normal pathways are congested or unavailable. Rerouting is used in the transportation sector (e.g., railroads), communications, and other sectors, in addition to the power systems.
- Substitution: Replacing a component or operating procedure with an equivalent or comparable one can continue operations if shortages or attacks occur. Power production with alternative fuel sources (e.g., gas versus coal) is a common example in power systems; switching from wireless to wired communication may be an option to cope with a radio frequency (RF) jamming attack.
- Islanding: Islanding is a form of dynamic compartmentalization that utilities may use in response to a disruption and to prevent larger cascading failures.

- Moving target defense (MTD) and Frequency Hopping: MTD technologies attempt to change fundamental system attributes (e.g., IP addresses) in a coordinated manner that appears to be random to potential adversaries. The frequent changes are intended to confuse adversaries, ultimately preventing or delaying compromise of systems. MTDs are being researched to protect against physical and cyber threats. Frequency hopping is a related concept for changing RF wireless communication channels in a seemingly random manner to protect against jamming attacks.
- Conservation and Rationing: When systems are stressed or under attack, cutting off non-essential functions may enable continuation of essential functions. For example, when hospitals lose power, they will frequently cancel or delay non-essential procedures to ensure the power from backup generators can be used for critical, life-saving procedures.
- Deception Networks: This technology is a relatively new concept for defending against cyber attacks. Deception networks emulate real networks so well that attackers will explore and attack these decoys instead of exploiting the real targets. Furthermore, network defenders can observe attacker techniques and build defenses against them.
- Ingenuity: Though difficult to characterize and plan for human ingenuity can sometimes significantly affect the resilience of infrastructures.

Restorative Capacity

When a system is damaged or compromised, its restorative capacity is the extent to which the system can be repaired rapidly and efficiently. The restorative capacity can be considered the last line of defense because it may not be necessary if the other capacities are sufficiently effective; furthermore, activation of restorative features is generally more costly than activation of anticipative, absorptive, and adaptive features. The effects of restorative features are generally intended to be permanent and longer lasting than those from adaptive features.

Examples of resilience enhancement features that contribute to a system's restorative capacity include:

- Graceful Degradation and Fail-safe Modes: When system operators have advanced notice of threats, they may switch to fail-safe operating modes and elect to gracefully degrade to prevent significant losses from an unplanned shut down. For example, petrochemical refineries in the Gulf Coast region often proactively shut down operations 48 to 72 hours before a hurricane makes landfall in the region. Doing so protects equipment from further damage from an unplanned shutdown; furthermore, the cost of not operating for 1 week is far less than the cost of additional equipment repairs and not operating for weeks or months while repairs are made.
- Intrusion Protection System (IPS): In cyber networks, IPSs can automatically implement defensive measures, such as closing access points and reconfiguring firewalls. Significant research is ongoing to develop autonomous methods that learn to detect and mitigate automatically cyber threats. IPS typically operates only when an IDS is used in combination with it.
- Fault Detection and Forensics: Fault detection technologies can notify operators about issues and enable a rapid response to repair them. Forensics, especially in cyber applications, are needed to identify the cause of failure or attack so that appropriate steps can be taken to repair the systems and prevent against future attacks.
- Reciprocal Aid Agreements: Utilities and industries will often participate in reciprocal aid agreements so that if one member is under duress, other members will share staff and/or equipment to limit damages and benefit the overall group. Electric power utilities,

telecommunications providers, and emergency services (e.g., fire departments) often exercise these agreements.

Considerations for Resilient Design

The resilience capacities and enhancement features can be used in a variety of manners. They can be incorporated into the design of new systems or they can be used to improve and address resilience deficiencies in existing systems. However, no singular combination of resilience enhancement features is optimal across all systems. Rather, operators and designers need to determine which features best achieve system operating goals, address threats of concern, and meet budget, regulatory, and other goals. This section discusses the many issues that must often be considered when making resilient design decisions.

System of Interest

Though saying so may seem trivial, the first factor that ought to be considered in the resilient design process is the infrastructure system being designed. Infrastructure systems can be large, complex systems with many functions and outputs. Trying to ensure that every element of an infrastructure system is protected and resilient can seem, if not be, impossible. Hence, when making design decisions, understanding of the critical infrastructure elements that enable completion of the most important missions is key. Designers need to address the following mission questions:

- What is the infrastructure system's mission(s)? Oftentimes, infrastructure mission can be described in terms of goods and services that the system provides.
- Of these missions, which are most critical and of highest priority? These missions should be the primary beneficiaries of resilience enhancements, and lower priority missions may need to be excluded from analyses if resources do not permit addressing every mission element. Prioritization requires analysts' preferences be included. For electrical power systems, priority may be given to loads supporting hospitals, police stations, and other emergency services. For the gas and oil industry, customers will sometimes opt for "interruptible" contracts that specify in times of shortages and crises, these customers are of the lowest priority and may be cut off, if necessary. In some commercial industries, the highest priority may be given to the goods that bring the company the highest revenue.
- How does the infrastructure system achieve its mission? Understanding the manner in which mission is achieved will identify critical components, devices, and processes that should be the focus of design improvements.

Another consideration is to determine the various time scales that exist within the system.

- How soon will negative consequences be realized if the infrastructure mission is compromised? Power failures are realized almost immediately, possibly resulting in rapid economic, security, and health impacts if the failures are not quickly remediated (within minutes to hours). Financial systems often move at rapid timescales; fractions of a second can make large monetary differences for stock trading companies. Petroleum transmission systems operate at a slower timescale. Reserves in storage facilities (and in the transmission pipes themselves) permit for disruption recoveries to occur at a slower pace, with lesser impacts to downstream customers. As a thought experiment, try to recall the last time a black out occurred and affected your day. How much did your day differ from how you expected it would go? Now, try to remember the

last time when a petroleum outage or shortage affected your day. Was your day significantly affected, aside from paying a few extra dollars at the gas pump?

- How long will the recovery take, and how do operations vary throughout the recovery duration? Extended recoveries may require rationing of resources to make sure they last for the entirety of the recovery. Additionally, activities may vary from one stage of the recovery to the next. For example, when a hurricane affects transportation networks, response and recovery begin with rerouting of shipments away from the affected area. After weather conditions stabilize, rerouting will likely continue while debris removal operations initiate. Transportation will only commence in the affected region after debris removal is sufficiently advanced and fuel is available.

Other timescale examples include:

- During a natural disaster, a refinery may be required to provide fuel throughout the long-time scale of recovering from the disaster. Although anticipative and absorptive features can provide immediate benefits, features contributing to the adaptive and restorative capacities will likely be needed to continue meeting goals throughout the extended recovery period.
- A nuclear reactor has hard, real-time communications requirements for safety controls. Rerouting communications in response to a cyber attack in a manner that adds too much latency to this time critical system would not be an acceptable option. An IDS that passively monitors the controls and alerts to anomalies might be acceptable if it does not add latency to the critical communication paths

A last consideration in identifying the system is to consider the degree of automation within the system. Degree of automation is closely related to timescales within the system. Highly automated systems may enable a faster response; however, they may introduce additional vulnerabilities. Attackers may know or learn the effects of automated responses and use them to their benefit by creating false positives. Conversely, automation may be required to attain necessary response speeds to address coordinated cyber attacks. An IDS may be effective for alerting to the presence of a threat. However, if threat mitigation is not automated and relies on a human response, then the response and recovery may not be on a fast-enough timescale to prevent substantial damage. Cyber attacks are a common example of this temporal asymmetry between attackers and defenders. Attacks may be automated, but responses to cyber attacks in the electrical power sector and other industries typically rely on predetermined playbooks in which decisions must be made by humans that move at “human speed.” A sophisticated attacker with a highly automated approach can often accomplish his/her goal well before the human response takes effect.

Threat Space

Resilience is a contextual concept and must be considered in the context of a specific disruptions and threats. A system may be highly resilient to one set of threats but vulnerable to other threats, so resilient design activities should identify the threats of greatest concern before proposing resilience enhancing features. Oftentimes, the number of scenarios postulated can seem endless, and constraints (time, resources, etc.) prevent consideration of all scenarios. In these cases, key stakeholders ought to be consulted to identify and prioritize which threat scenarios should be included in design efforts. With

the selection of threat scenarios, designers can then match the mitigations that best address those scenarios.

Development of the scenarios should include specification of the disruption, the effect that the disruption has on the system, timing of effects, and system response mechanisms—planned or implemented. Common disruptions considered in resilience analyses consist of natural disasters, accidents, and malevolent events. Examples of threat space considerations and how they affect the choice of resilience enhancing features include:

- Buildings concerned about power loss during storms and flooding should not house backup generators and related equipment in basements or floors that could experience flooding.
- HAZMAT suits and protective equipment necessary to respond to dangerous chemical explosions should not be located in the path of expected chemical plumes.
- Information technology networks concerned about insider threats should not rely solely on firewalls preventing unauthorized, external access. These networks likely require additional network analytics that can recognize anomalous behaviors by users with credentialed access.

Operational Constraints

Resilient design activities often focus on the effect that proposed mitigations have upon reducing negative consequences from threats. However, in practice, designers need to consider the effects of the mitigations on other operational considerations. Common considerations include:

- Budget and Cost. What is the cost for the proposed features? Cost estimates should include both upfront investment and the cost to operate, maintain, and implement when threats are realized. These costs need to be weighed against available budgets, the potential benefits from the mitigation (which may include loss avoidance), decreased insurance premiums, competitive advantage gained from decreased down time, etc. For example, building a fully redundant system may increase resilience; however, doing so is generally cost prohibitive.
- Reliability. Electrical power utilities (and other industries) are regulated according to rigid reliability standards. Design modifications intended to enhance resilience that would negatively affect reliability of power systems would likely be rejected immediately.
- Safety. Safety is another common requirement (and regulatory basis) that must be satisfied. Safety is the primary requirement of nuclear power and chemical manufacturing facilities, so resilient design features must be proven to ensure continued compliance with safety requirements.
- Other regulations. Beyond reliability and safety, many industries must comply with additional regulations. Banking and healthcare industries have privacy standards that they must maintain, so the proposed cyber-resilience measures must not expose client information. Food preparation companies and supply chains must meet health standards; therefore, rerouting efforts that keep foods in transit for periods of time that violate health standards would not be effective.
- Size, weight, and power (SWaP). Cars, planes, satellites, and other vehicles have stringent SWaP restrictions that resilient design efforts must consider. Inclusion of sophisticated, computationally intense IDS in satellites with limited battery capacities would likely be rejected due to SWaP constraints.

Summary

Four resilience capacities (anticipative, absorptive, adaptive, and restorative) contribute to an infrastructure's resilience. Resilience enhancement features are the tangible technologies, designs, and procedures that are implemented, contribute to one or more of the capacities, and determine the overall resilience of the system.

Design of resilient systems does not include a "one-size-fits-all" approach. Selection of the right resilience enhancement features requires consideration of several factors. Designers need to consider the infrastructure system's core missions, the potential threat space, and a variety of operational constraints. Development of formal, resilient design processes is a subject of continuing research, and this section detailed resilient design considerations.

References

[1] Vugrin et al., 2011.

Biringer, B., E. Vugrin, and D. Warren, *Critical Infrastructure System Security and Resiliency*, CRC Press: Boca Raton, Florida, 2013.

[2] Dragos, Inc., 2017a, "CRASHOVERRIDE Analysis of the Threat to Electric Grid Operations." Technical Report Version 2.20170613, Dragos Inc., Hanover, MD, 2017.

[3] Dragos, Inc. 2017b, "TRISIS Malware Analysis of Safety System Targeted Malware." Technical Report version 1.20171213. Dragos Inc., Hanover, Maryland, 2017.

[4] Netflix, "Github – Netflix/chaosmonkey: Chaos Monkey is a resiliency tool," 2019, Accessed December 9, 2019, at https://github.com/Netflix/chaosmonkey

[5] Ross, R., V. Pillitteri, R. Graubart, D. Bodeau, R. McQuaid, *Developing Cyber Resilient Systems: A Systems Security Engineering Approach*, National Institute of Standards and Technologies, NIST SP 800-160, Volume 2, Gaithersburg, Maryland, 2019.

Thoughtful Questions

1. Give an example of a resilience enhancement feature that can be viewed as contributing to more than one resilience capacity.
2. Resilience enhancing features may be contradictory (i.e., they may enhance one capacity while degrading another). Give an example of such a feature.
3. It may not be financially possible to invest in all resilience enhancing features. Describe a method you would use to choose a subset of features that you would propose to an infrastructure manager.
4. Cascading failures between system components is an important resilience consideration when defining the system of interest. Give an example of an infrastructure with the potential for cascading failures. What resilience enhancing features could benefit your example?
5. Give an example of two contrasting infrastructure systems: one that needs to obtain a certain level of performance at all costs and one that can tolerate degraded performance in order to conserve resources.
6. Give an example of a system where humans should be considered part of the infrastructure system itself. How does this affect which resilience enhancing features may be considered?
7. Systems can be studied at a variety of different scales. Suppose the threats and proposed resilience enhancing features to the system have already been chosen, what are some questions one could ask to determine the proper level of system granularity to consider for the analysis?
8. An important consideration when considering resilience enhancing features is identifying when the response/recover is deemed to be complete. Give an example identifying these criteria for a system under a specific disruption.
9. Many experiments and analyses include sources of uncertainty. Give some examples of uncertainty as they relate to resilience enhancing features discussed in this chapter.

Further Reading

- Biringer, B., E. Vugrin, and D. Warren, Discussion and application of resilience capacities: Chapter 10 of *Critical Infrastructure System Security and Resiliency*, CRC Press: Boca Raton, Florida, 2013.
- The White House, *Presidential Policy Directive 21 (PPD-21) Critical Infrastructure Security and Resilience*, February 12, 2013.
- Ross, R., V. Pillitteri, R. Graubart, D. Bodeau, R. McQuaid, Cyber Resilience Framework: *NIST SP 800-160, Volume 2, Developing Cyber Resilient Systems: A Systems Security Engineering Approach*.
- Rieger, C. G., D. I. Gertman, M. A. McQueen, Resilient Control Systems: Next Generation Design Research, *2nd IEEE Conference on Human System Interaction*, Catania, Italy, May 2009.