ESTCP

FINAL REPORT

Demonstration of a Novel Technology to Manage Electricity Demand in Grid-Independent Military Microgrids

Korbaga Woldekidan, Donald Price, Tobias Shapinsky, Maggie Sullivan, and Eric Bonnema
National Renewable Energy Laboratory

Peter Curtiss, Scott Roberts, Chris Rowland, and John Gehring
S&C Electric Company

September 2025

ESTCP Project: EW22-7512

This work was authored by NREL for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308.. Funding provided by the Department of Defense's Environmental Security Technology Certification Program (ESTCP) under Agreement IAG-21-18253-8. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE	**2. REPORT TYPE**	**3. DATES COVERED**
07-31-2023	ESTCP Technical Report	May 2022 – Dec 2024

4. TITLE AND SUBTITLE

Demonstration of a Novel Technology To Manage Electricity Demand in Grid-Independent Military Microgrids

5a. CONTRACT NUMBER
IAG-21-18253-8

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Korbaga Woldekidan, Donald Price, Tobias Shapinsky, Maggie Sullivan, and Eric Bonnema, Peter Curtiss, Scott Roberts, Chris Rowland, and John Gehring

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

National Renewable Energy Laboratory, 15013 Denver W Pkwy,
Golden, CO 80401
S&C Electric Company, 6601 N Ridge Blvd, Chicago, IL 60626-3997

8. PERFORMING ORGANIZATION REPORT

NREL/TP-5500-90715

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

The Environmental Security Technology Certification Program office
4800 Mark Center Drive, Suite 16F16, Alexandria, VA 22350-3605

10. SPONSOR/MONITOR'S ACRONYM(S)
ESTCP

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
EW22-7512

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

When military microgrids isolate themselves from the main electrical grid, they must balance electricity supply and demand locally. Since local generation may be limited, the current strategy is to shed all but the most critical loads by turning off smart circuit breakers. This strategy is typically applied at the building level, meaning that entire buildings housing mission-critical activities must be excluded from any load management, while those deemed noncritical may lose service entirely.

This report demonstrates the use of automated demand response (ADR) technology to facilitate communication between a microgrid controller (MGC) and building management systems (BMS) to control the loads of the microgrid in islanded mode. It also describes the software updates that need to be implemented by the MGC to support this communication. This automated approach achieves load shedding and load shifting through communication signals sent to equipment controllers, rather than by cutting off the flow of electricity within the microgrid itself. Because this strategy works only within the base network and without connection to external entities, it avoids the major cybersecurity issue that has been encountered with previous applications of ADR on military bases.

15. SUBJECT TERMS
Automated Demand Response, Microgrid, Microgrid controller, Load Shifting, Load Shedding

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
					Korbaga Woldekidan
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	UNCLASS	81	**19b. TELEPHONE NUMBER** *(include area code)*
UNCLASS	UNCLASS	UNCLASS			315-664-0594

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

FINAL REPORT
Project: EW22-7512

TABLE OF CONTENTS

Page

LIST OF FIGURES

Page

LIST OF FIGURES

Page

LIST OF TABLES

Page

ACRONYMS AND ABBREVIATIONS

ADR	automated demand response
ATO	authority to operate
BMS	building management system
DOD	U.S. Department of Defense
DOE	U.S. Department of Energy
ESIF	Energy Systems Integration Facility
FUT	facility under test
HMI	human-machine interface
HTTP	hypertext transfer protocol
HTTPS	hypertext transfer protocol secure
HVAC	heating, ventilating, and air conditioning
IP	Internet Protocol
JB MDL	Joint Base McGuire-Dix-Lakehurst
kW	kilowatt
LAN	local area network
MGC	microgrid controller
NIC	network interface card
Nmap	Network Mapper
NREL	National Renewable Energy Laboratory
OpenADR	Open Automated Demand Response
SCADA	supervisory control and data acquisition
TLS	Transport Layer Security
VEN	virtual end node
VPN	virtual private network
VTN	virtual top node

KEYWORDS

automated demand response, building management system, island mode, load shedding, microgrid, microgrid controller, OpenADR

ABSTRACT

INTRODUCTION AND OBJECTIVES

This project, "Demonstration of a Novel Technology To Manage Electricity Demand in Grid-Independent Military Microgrids," was proposed and approved to demonstrate the use of automated demand response (ADR) technologies to manage building loads within a microgrid. These technologies may be used to maintain mission-critical functions and microgrid stability when microgrids are isolated from the main electrical grid and to enable smoother transitions between grid-connected and grid-independent modes.

TECHNOLOGY DESCRIPTION

Upon disconnecting from the main electrical grid, an installation's microgrid controller (MGC) will evaluate the balance between the supply and demand of power. If the controller anticipates a supply shortfall, it sends a load shedding request to a local Open Automated Demand Response (OpenADR) server via the base's secure communication network to maintain safe operating voltages. The OpenADR server signals the installation's building management systems (BMS) to initiate load shedding control sequences, according to the requested level of load shedding. Mission-critical equipment can opt out of any demand response. Once reconnected to the grid, the MGC returns the installation to normal operation via the same OpenADR client-server architecture.

RESULTS

The National Renewable Energy Laboratory (NREL), in collaboration with S&C Electric Inc., conducted laboratory and field tests of the proposed technology. Both tests successfully demonstrated seamless communication between the MGC and the BMS via the OpenADR server. Although the solution demonstrates effective communication, latency considerations within the round-trip communication path require careful application. This method is well suited for load shedding/restoration during transitions between islanded and grid-tied modes, for managing normal load fluctuations, and for planned maintenance scenarios. However, due to the inherent latency, this method is not recommended for scenarios that require rapid load shedding (within 1 or 2 seconds), such as sudden load increases or unexpected power source failures.

BENEFITS

This technology enables substantial load shedding across an entire installation without having to shut down entire buildings. By using a standardized, open communication protocol (OpenADR) operating only on the base's network, the proposed approach is secure against cyber threats, requires little customization, and works with a wide range of MGCs and BMS. The use of BMS routines to shed and restore load in a controlled manner will reduce the wear and tear on equipment caused by the abrupt power changes that characterize the transition between operating modes in today's microgrids. Intelligent load management avoids large power spikes caused by coincident startup loads and allows engineers to reduce the capacity of a microgrid's generation assets, backup storage, and distribution infrastructure, further reducing capital costs.

EXECUTIVE SUMMARY

INTRODUCTION

This research was conducted by the National Renewable Energy Laboratory (NREL) in collaboration with S&C Electric Inc. with funding from the U.S. Department of Defense's Environmental Security Technology Certification Program. The project demonstrates the use of cybersecure automated demand response (ADR) technology to effectively manage microgrid loads during grid-independent, or "islanded," operation.

When military microgrids are isolated from the main electrical grid, they must balance electricity supply and demand locally. Because local generation may be limited, the predominant strategy is to switch off all but the most critical loads by tripping smart circuit breakers, which must then be manually reset. This approach is usually implemented at the building level, which means that the buildings with mission-critical activities are exempt from load management and remain fully powered, while the buildings classified as noncritical may be affected by a complete outage.

In this research, we used a local ADR server to facilitate communication between a microgrid controller (MGC) and a building automation system. The building automation system selectively controls its assets in response to the load shedding requests from the MGC, avoiding a total outage as opposed to the conventional control approach. Because it operates only on the base network, with no connection to external entities, this strategy avoids the main cybersecurity concern raised by past applications of ADR on military bases.

OBJECTIVES

The objective of this project is to demonstrate the use of ADR technologies to manage building loads within a microgrid for the purposes of maintaining mission-critical functions and microgrid stability when isolated from the main electrical grid and enabling smoother transitions between grid-connected and grid-independent modes. Specific technical objectives and outcomes of the proposed demonstration include:

- Demonstrate multi-building demand management in concert with the wider microgrid control strategies.
- Provide methodology for facility managers to maintain control of building load shedding.
- Validate the efficacy of greater granularity of load shedding.
- Demonstrate measurable demand reduction.
- Validate reduced installation and maintenance costs (e.g., of additional switchgear).
- Demonstrate scalability of the load management approach across multiple buildings.

TECHNOLOGY APPROACH

The technology employs a specific ADR communication protocol and architecture called Open Automated Demand Response (or OpenADR). It is a standardized, interoperable, and secure protocol for transmitting demand response signals, originally developed to help electric utilities manage periods of critical peak load. These signals are transmitted from a server (a "virtual top node," or VTN in OpenADR parlance) to clients (also known as "virtual end nodes," or VENs), which communicate with end-use equipment controls (including building management systems [BMS]) to automatically change their operation, in order to shed or shift building loads (e.g., heating, ventilation, air conditioning, lighting, water pumping, and other miscellaneous motors). It is important to note that an ADR server located outside of a U.S. Department of Defense (DOD) microgrid can only modulate loads inside the microgrid during routine, grid-connected operation. This approach also currently poses potential cybersecurity concerns because of its reliance on connection to external networks.

In contrast, this technology makes OpenADR much more valuable to DOD by placing an ADR server node (the "Local Automated Demand Response Server" in Figure ES.1) behind the installation's network firewall (i.e., with no external network connection) and tightly coupling it with a DOD MGC via a generic communications protocol. By doing so, a lightweight, interoperable communications layer is created to deliver balancing services during grid-independent episodes, when extreme events or other power disruptions have taken down the main grid. In this case, the MGC will instruct the OpenADR VTN to issue commands across the secure microgrid communications network to OpenADR VENs to trigger the appropriate control modes in end-use equipment to help balance the supply and demand of electricity. Rather than an all-or-nothing approach, which preserves critical base functions while shedding large blocks of demand using smart breakers (the current state of the art in military microgrids), this ADR-based approach would achieve greater resilience by incorporating the expertise of building managers regarding the priority of different loads across the installation at any given time.

The MGC's algorithms are updated to consider the building loads in the wider microgrid strategy (i.e., for demand reduction). By integrating the MGC with a commercially available OpenADR solution, this project introduced a scalable solution for islanded microgrid load control. As long as an installation's MGC and BMS are OpenADR-compliant, the proposed solution is vendor agnostic.

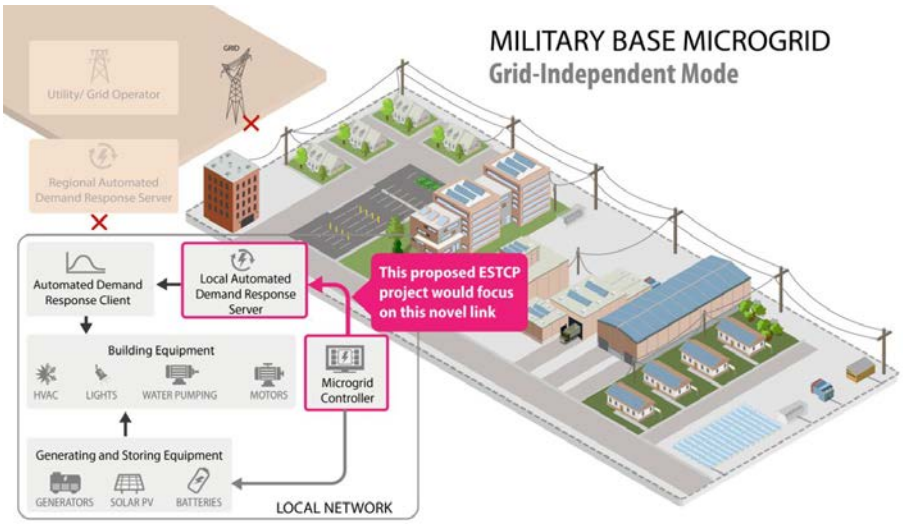

Figure ES.1 Novel use of a local ADR server to manage loads in a grid-independent microgrid

RESULTS AND DISCUSSION

Both laboratory and field tests were conducted, considering the load shedding required to achieve a closed transition to islanded operation and to account for changing loads and the availability of power sources while islanded. For these tests, a commercial MGC was used with normal programming within the expected constraints of system capabilities. This means the solution did not require any specialized modification to the code base of the controller.

Considering the latency of the round-trip communication path between the MGC and the different devices/assets involved in the load shedding processes, there are certain scenarios for which the demonstrated solutions are suitable and others for which they are not. When islanded, the method described in this report could be used to set target power demands for buildings within the microgrid. This would be done at the time of islanding and be based on the mission criticality of each building. A custom process would be created that compares the sum of capacities of all available power sources to the sum of all anticipated building loads. If there is insufficient generation, then OpenADR would be used to limit/shed building loads to manageable values. The method can also accommodate normal variations in load and the need to remove a power source from operation for maintenance.

Due to the inherent latency of system response to a shedding request, the developed methodology cannot be used in situations that require fast response, such as:

- Decreasing loads rapidly as necessary due to an unexpected loss of a power source.
- Attempting a black start that requires pre-shedding of loads prior to energizing the bus.
- Reducing demand on a microgrid due to a sudden and unexpected increase in load.

An assessment of the methodology's vulnerability to cybersecurity attacks revealed the need for robust security hardening. This report outlines cybersecurity recommendations essential for safe operation of the technology.

IMPLICATIONS FOR FUTURE RESEARCH AND BENEFITS

The ADR framework will eliminate the need for expensive, remotely controlled switchgear, reducing microgrid capital and maintenance costs. It will also enable substantial load shedding across an entire installation without entire buildings having to go dark. By using a standardized, open communication protocol (OpenADR) operating only on the base's network, the proposed approach will be secure against cyber threats, require little customization to deploy, and will work with a wide range of MGCs and BMS. Using BMS routines to shed and restore load in a controlled manner will reduce wear and tear on equipment caused by the abrupt power changes that characterize the transition between operating modes in microgrids today. Intelligent load management will avoid large power spikes from coincident startup loads, allowing engineers to reduce the capacity of a microgrid's generation assets, backup storage, and distribution infrastructure, further lowering capital costs.

In short, the proposed approach to managing microgrid loads will allow the DOD to leapfrog to an agile, scalable, and secure technology platform for preserving mission-critical installation functionality during adverse conditions. By implementing OpenADR to mediate communications between MGCs and building BMS, the DOD will leverage a broad trend in the building control industry toward improved interoperability and increased adoption of intelligent flexible load management, resulting in reduced costs and streamlined deployment. If OpenADR were to achieve an authority to operate (ATO) at some point in the future for utility ADR participation, it could potentially be used to manage installation loads during both grid-connected and islanded operating modalities, with seamless toggling between the two.

Benefits to DOD include:

- Elimination of addition of remotely operated switchgear (at $30,000–$50,000 per building) for more granular control.
- Reduced risk of damage to equipment from abrupt interruption of electric service.
- Feasible much longer duration in grid-independent modes, with continued use of non-mission-critical buildings (albeit potentially with reduced lighting and less comfortable temperatures).
- Ability to take advantage of the fast private sector advancements in ADR technologies, while also avoiding one-off overly customized solutions that are costly, fragile, and have limited support.
- Unnecessary loads can be curtailed automatically, while also enabling nimble human-in-the-loop actions.
- Ability to explore "what-if" scenarios, to provide operational guidance on load management needed to preserve critical base functions under various extreme circumstances an installation could potentially encounter. Other military analytics for battlefield circumstances offers "what-if" scenario analysis in near real time to provide tactical guidance. Our energy system would provide analogous services.

1.0 OBJECTIVE

The objective of this project is to demonstrate the use of automated demand response (ADR) technologies to manage building loads within a microgrid for the purposes of maintaining mission-critical functions and microgrid stability when isolated from the main electrical grid and enabling smoother transitions between grid-connected and grid-independent modes.

Laboratory testing and field testing were used to demonstrate the ability of a cybersecure ADR server-client architecture to manage microgrid building loads for the purposes of maintaining critical base functions during islanding and enabling smoother transitions between grid-connected and islanded modes. Technical objectives and outcomes of the proposed demonstration include:

- Demonstrate multi-building demand management in concert with the wider microgrid control strategies.
- Provide methodology for facility managers to maintain control of building load shedding.
- Validate the efficacy of greater granularity of load shedding.
- Demonstrate measurable demand reduction.
- Validate reduced installation and maintenance costs (e.g., of additional switchgear).
- Demonstrate scalability of the load management approach across multiple buildings.

2.0 BACKGROUND

When military microgrids are isolated from the main electrical grid, they must balance electricity supply and demand locally. Because local generation may be limited, the prevailing strategy is to shut down all but the most critical loads by tripping smart circuit breakers, which must then be manually reset (e.g., U.S. Department of Defense [DOD] Environmental Security Technology Certification Program project EW-201350). This approach is generally implemented at the building level, which means the buildings with mission-critical activities are exempt from load management and remain fully powered, whereas those deemed noncritical can experience a complete loss of service. The remotely controlled switchgear essential for managing loads in this manner is costly, ranging from $30,000 to $50,000 per building. Although this strategy is effective for load shedding, it can disrupt operational continuity and potentially damage equipment during both the disconnection and reenergization process.

In this project, we demonstrated the implementation of cybersecure ADR technology to effectively manage microgrid loads during grid-independent, also known as "islanded," operation. Our objectives included reducing costs, safeguarding equipment, and improving the resiliency of DOD microgrids. This innovative method facilitates load shedding and shifting via communication signals to equipment controllers, rather than interrupting the electrical flow within the microgrid. As the system operates exclusively on the base network and does not connect to external parties, it mitigates the primary cybersecurity risks identified in previous ADR applications on military installations.

The efficacy of the technology was proven both in a controlled laboratory environment and in an operational DOD facility. The laboratory test was conducted at the National Renewable Energy Laboratory's (NREL's) Energy Systems Integration Facility (ESIF) in Golden, Colorado. A hardware-in-the-loop setup using the Alfalfa (Benne et al. 2018) simulation platform to emulate the effects of actual controller decisions on building energy consumption and performance was used for the laboratory test. The field test was conducted at a pilot building located at the Joint Base McGuire-Dix-Lakehurst (JB MDL) in McGuire Air Force Base, New Jersey. The test demonstrated successful use of the technology to selectively manage building assets. This capability enabled us to fulfill load shedding requests from a microgrid controller (MGC) during real-world conditions.

3.0 TECHNOLOGY DESCRIPTION

This project employs ADR technology to facilitate communication between an MGC and a building management system (BMS). Many ADR solutions are built on the OpenADR communications protocol, which originated in 2002 (OpenADR Alliance 2012). To support this development, the OpenADR Alliance was established in 2010 by a consortium of utilities, software vendors, device manufacturers, national laboratories, demand response aggregators, testing and certification labs, system integrators, and consulting firms. The alliance currently boasts more than 200 members and has been instrumental in the creation and adoption of OpenADR—more than 280 commercial products are now compliant with its specifications (OpenADR Alliance 2023).

In this project, we have integrated commercial server and client ADR products from GridFabric, a company whose founders not only developed the initial OpenADR software but also have played a pivotal role in its continued advancement. OpenADR empowers devices and systems to receive signals that prompt them to adjust electricity usage, either by reducing consumption or shifting it temporally.

The core components of OpenADR communication are the virtual top node (VTN) and virtual end node (VEN). The VTN, often managed by a utility or an aggregator, sends out demand response directives to the VEN, which represents the end-user equipment, such as smart thermostats, energy management systems, or controllable appliances. VENs autonomously alter their energy use in response to these signals, adhering to preestablished strategies or user settings. They also provide feedback to the VTN, allowing for the monitoring and coordination of the demand response event.

Using the OpenADR protocol, the VEN registers with the VTN, retrieves demand response events, reacts accordingly, and relays system performance and status information back to the VTN. This two-way communication ensures demand response is both effective and efficient, optimizing energy consumption while maintaining user comfort and operational requirements.

OpenADR distinguishes itself from conventional building control protocols such as Modbus and BACnet in several key aspects. First, it employs hypertext transfer protocol secure (HTTPS) as its primary communication method, establishing a secure channel for the transmission of OpenADR messages. This protocol is the global standard for securing web traffic. Second, OpenADR communications are typically scheduled to occur once per minute. This frequency is sufficient for most OpenADR use cases, which involve demand response events.

In this project, NREL and S&C Electric Inc. have implemented an innovative and advanced load management strategy for islanded microgrids, leveraging ADR to facilitate communications between an MGC and a BMS. This approach allows for the reduction of demand across facilities without necessitating a complete shutdown of building operations. Even with significant demand reduction, building occupants can maintain productivity, albeit potentially with reduced comfort. Building operators can use the familiar BMS human-machine interface (HMI) to opt out of demand reduction for specific loads or the entire building during critical missions.

As shown in Figure 1, the MGC monitors both generation capacity and various assets within the microgrid and issues a demand reduction request when necessary. This request is communicated to the ADR server (VTN) via Modbus. Upon receiving the request, the VTN forwards it to the

BMS (VEN) using the OpenADR protocol. For this project, we have used a Tridium BMS, but the methodology also could be used by other BMS. The BMS is preconfigured with various levels of load reduction strategies—low, medium, and high—based on the controllable assets available, allowing for a nuanced approach to demand reduction.

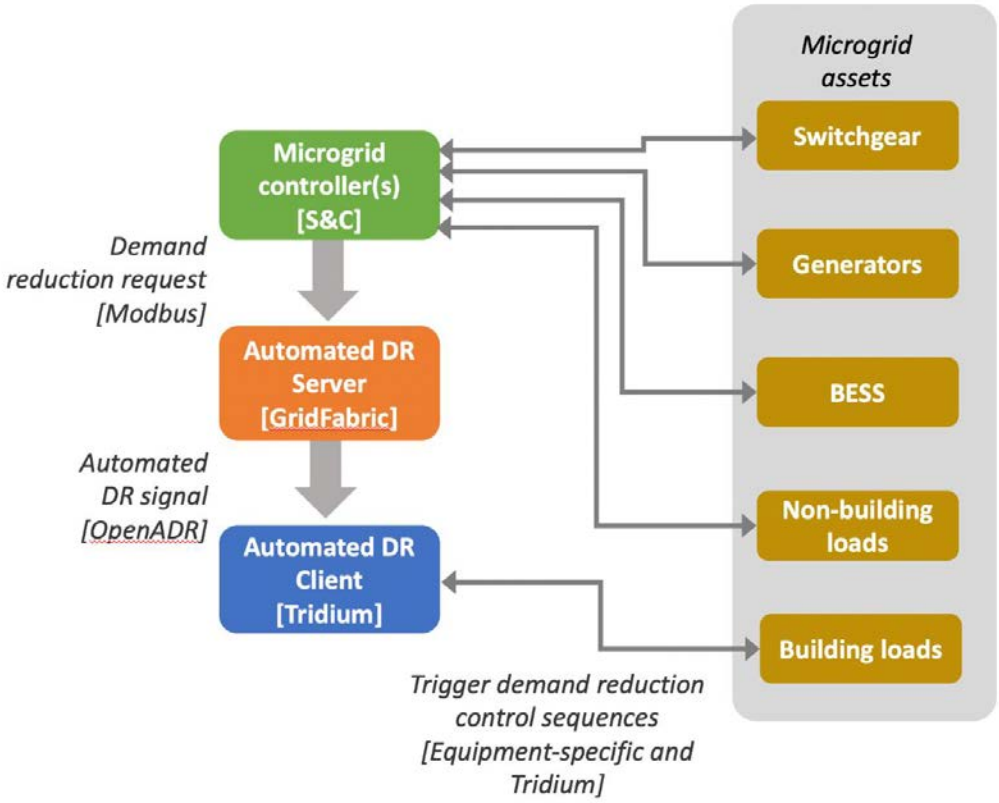

Figure 1. Novel use of a local ADR server to manage loads in a grid-independent microgrid.

BESS = battery energy storage system; DR = demand response; S&C = S&C Electric Inc.

Figure 2 demonstrates the workflow for the load shedding request from the MGC. The MGC determines the maximum allowable load on the grid by considering the available generator capacity. It also sets the hysteresis margin that is used to determine when it is appropriate to restore loads. These parameters are then communicated to the BMS via the OpenADR driver.

If the building load exceeds the maximum allowable load, the BMS initiates the necessary load shedding measures until the campus load is reduced to a level below the maximum allowable load. In cases in which the available load shedding capacity is insufficient to fulfill the load shedding request, the MGC resorts to the conventional control methods to manage the system, such as disconnecting the noncritical buildings from the grid.

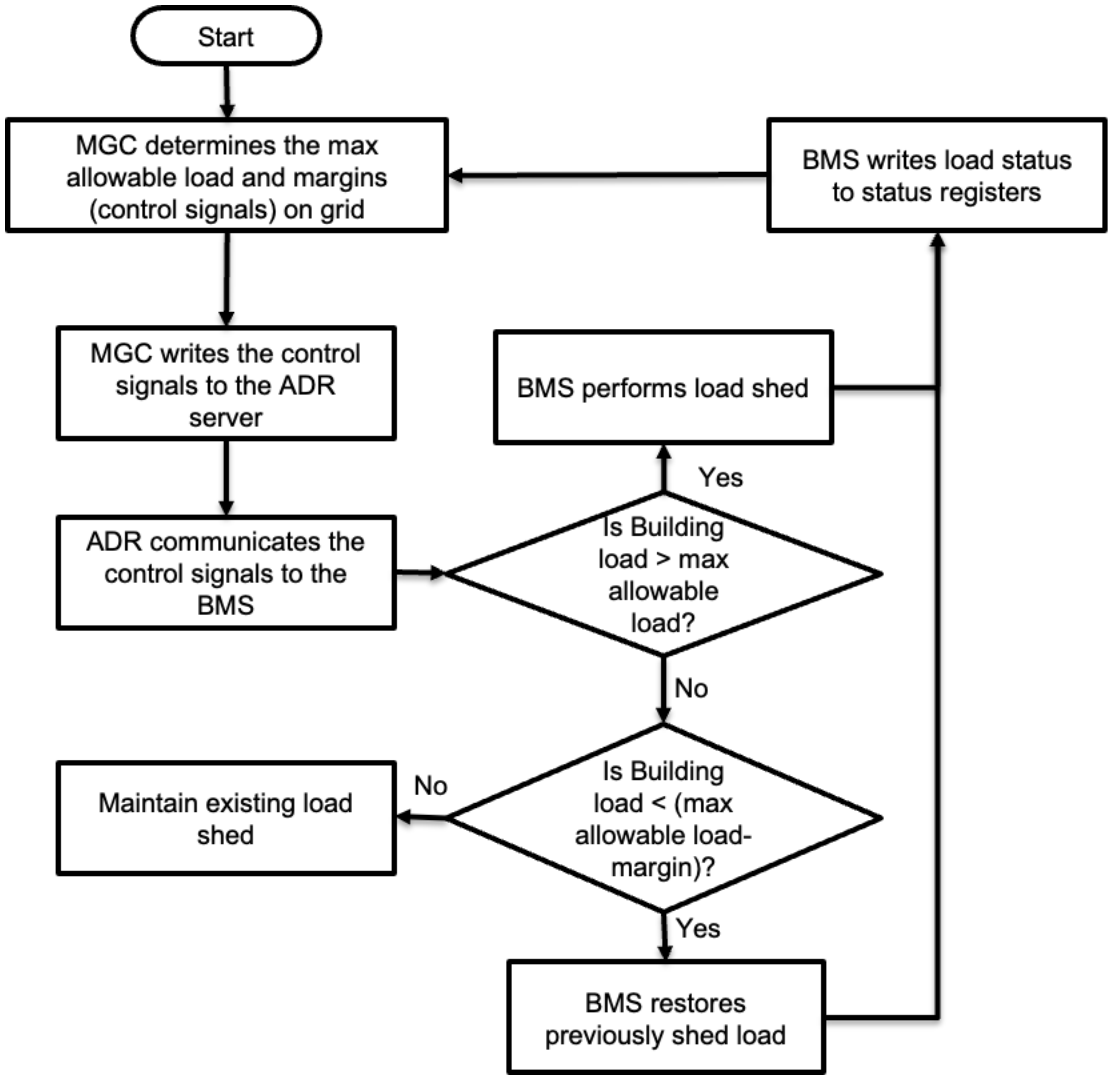

Figure 2. Interaction between MGC, ADR Server, and BMS for load power threshold set point control

The following sections discuss how communication between different components is established.

3.1 INTEGRATION BETWEEN MGC AND OpenADR SERVER

One of the novel contributions of this project is linking an MGC with ADR technologies. There is preexisting commercial software to facilitate OpenADR communication between the Tridium BMS and the VTN. However, at the start of this project, there was no standard solution for connecting the S&C MGC to an OpenADR VTN. Therefore, we had to come up with a solution to bridge the capabilities of these two systems.

We settled on an architecture with a piece of middleware between the MGC and the VTN that would allow for bidirectional communication. This takes the form of an application that provides Modbus registers for the MGC to write to and periodically sends hypertext transfer protocol

(HTTP) requests to the VTN to create and update events as well as to retrieve report data. This application will be referred as the "translator" throughout this paper.

The communication works as follows: The translator makes periodic calls to the VTN to retrieve the current values for the "Load power," "Available load shed," and "Baseline load," which it then writes to the Modbus registers, corresponding to those values as laid out in Table 1. It then reads the "Upper allowed load threshold" and "Hysteresis margin" from the Modbus registers and constructs a request to send to the VTN to create or update a demand response event with those values.

Table 1. Modbus Registers for OpenADR Power Threshold Control Data Exchange

Reg	Name	Units	Type	Description	Sources
40010/40011	Upper allowed load threshold	Kilowatts (kW)	Float	Max load to be allowed by OpenADR in kW; above this value, loads need to be shed by OpenADR	Written by MGC; read by OpenADR
40020/40021	Hysteresis margin	kW	Float	Margin between the value in register 40010/40011 and the point at which loads are restored	Written by MGC; read by OpenADR
40030/40031	Load power	kW	Float	Present value of load in kW as measured by the OpenADR	Written by OpenADR; read by MGC
40040/40041	Available load shed	kW	Float	The estimated amount of load that can be shed as determined by the OpenADR	Written by OpenADR; read by MGC
40050/40051	Baseline load	kW	Float	Amount the OpenADR system would use if it was not in an event	Written by OpenADR; read by MGC

The translator alone is not enough to facilitate the communication. The VTN also needs to be modified, as it does not provide an out-of-the-box interface for writing events or retrieving report information. This is done by creating a custom adapter, which uses a programming framework provided by the VTN to expose end points that the translator queries to accomplish the given task.

3.2 INTEGRATION OF OpenADR CLIENTS WITH EQUIPMENT

OpenADR clients (the VEN) are responsible for responding to demand response events and requests from the OpenADR server (the VTN). In the real world, the VTN is controlled and run by the power grid utility operators. In this project, we repurposed it in such a way that the VTN runs locally on a GridFabric Canvas VTN docker instance. The VEN is integrated into the Tridium BMS as a driver module. The VEN driver allows direct communication to the VTN, enabling building control logic to respond to events quickly. The VTN and the VEN go through a handshake

process to establish communications. The VEN only needs the VTN URL and the VTN Id information to initiate the handshake. With the VTN information configured in the VEN, a Create Party Registration is initiated from the VEN. The VTN will receive the registration request, accept, and create the VEN in the VTN configuration. After the VEN requests the party registration again or goes through a forced registration from the VTN, the VEN will automatically be populated with the information to establish full communications, as shown in Figure 3.

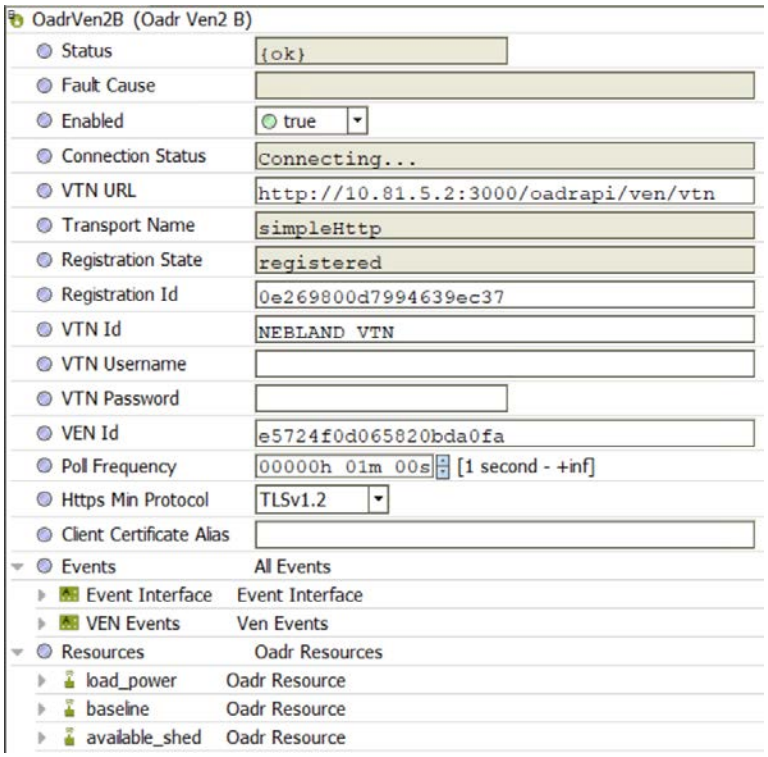

Figure 3. VEN registration in OpenADR driver module

The VEN sends reports to the VTN upon requests. As discussed in Section 3.1, the VEN receives the values of the "Upper allowed load threshold" and "Hysteresis margin" as VEN events from the VTN and sends back the values of "Load power," "Available load shed," and "Baseline load" to the VTN as a report. These values are then communicated to the MGC through the Translator.

Both the events and the reports are pulled from a VEN Resource. A VEN Resource is how information comes into and out of the VEN to the VTN. Resources are targeted sources of information, such as the load power meter of a building. If the VEN does not have the appropriate resource, the VEN will automatically opt out of an event or request from the VTN. The appropriate resource is linked to building control logic, to allow live changes in how the various equipment is controlled. For example, the "Upper allowed load threshold" tells the building control logic the maximum allowable power usage of the OpenADR-managed load that the grid can currently support. The building control logic will make an algorithmic decision on how to respond to the new maximum allowable power, shedding load as necessary.

4.0 RESULTS AND DISCUSSION

The proposed technology is demonstrated both in a lab setup and in the field. This section discusses the methodologies and findings.

4.1 LAB DEMONSTRATION

A hardware in-the-loop setup shown in Figure 4 is used for the lab demonstration. The setup includes a GridMaster MGC, a Tridium BMS controller loaded with the VEN driver, OpenADR server (VTN), a Modbus translator software, GridSim software, and a building model. The GridSim program is a power flow simulation of the grid that accounts for generator output, building loads, and switchgear operation. It simulates the generation capacity of the microgrid and opens/closes switchgears upon receiving commands from the MGC. It also simulates other loads connected to the microgrid such as Load 2 (building-2) and Load 3 (building-3), as shown in Figure 4. We considered three different buildings for the lab test. Building-1 (Load 1) represents a noncritical building that could be controlled through a BMS to provide load shedding based on a request from the MGC. Load 2 represents a critical building, and Load 3 represents a noncritical building, both controlled by a microgrid switchgear. Because Load 2 is a critical building, it will be disconnected from the grid only if Load 1 and Load 3 are not able to provide the required load shedding. If the load shedding request is within the load shedding limit of Load 1, the MGC will send the allowable load that has to be maintained through the BMS of the building. If the load shedding request is greater than the load shedding limit of Load 1, the MGC will disconnect Load 3, Load 1, and Load 2 sequentially until the load shedding request is met.

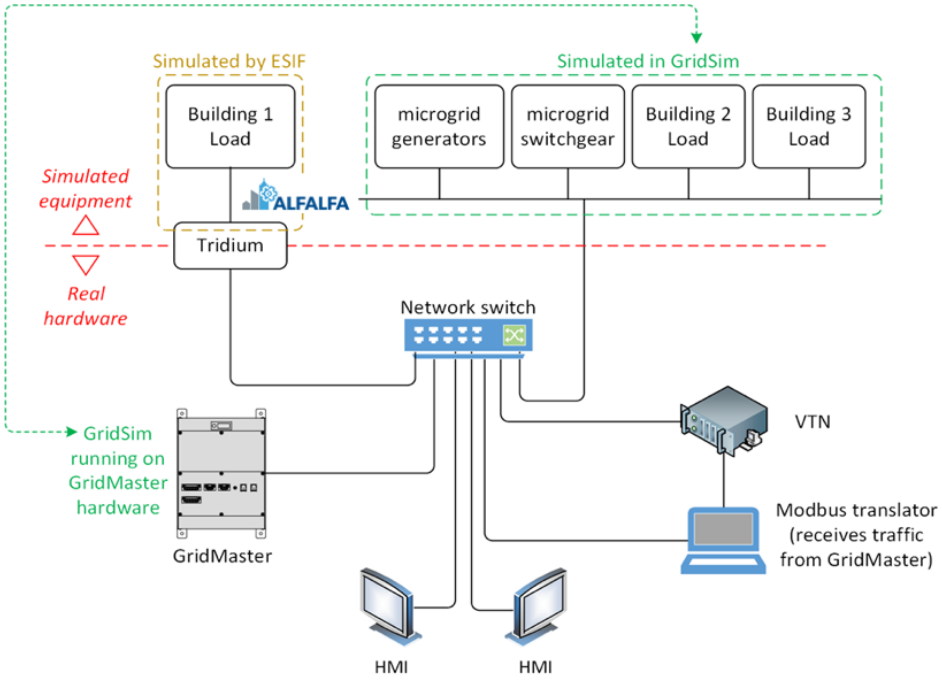

Figure 4. Laboratory test bed network topology.

ESIF = Energy Systems Integration Facility (National Renewable Energy Laboratory).

The communication between the different components in the setup are discussed in more detail in the following subsections.

4.1.1 Integration Between BMS and Buildings Simulation Model

To demonstrate the interaction between a BMS and buildings under its control in a lab setup, we used a hardware-in-the-loop approach, in which a building simulation model is connected to a Tridium BMS using a platform called Alfalfa. Alfalfa is a building energy modeling tool that builds on top of many years of energy modeling research. Traditionally, energy modeling tools run for a defined period before providing an output. Although this is good for many applications, it does not allow interactions with systems that need to run in real time. Alfalfa provides an interface that allows users to run models built from EnergyPlus in real time with interactivity.

In the lab setup, we used three building models interacting with a Tridium controller through Alfalfa. Alfalfa converts the necessary outputs from the simulation run such as temperature, power, set points, and others to BACnet points so that they can be easily discovered by controllers that support a BACnet communication protocol like the Tridium controller. The BACnet points are then actuated using a control logic developed in the controller. By doing this, we can prove that such a control strategy not only works in theory but also works when using real-world hardware and control protocols. Discovering and resolving issues in the lab allows us to conduct more efficient field tests.

The following subsections discuss details of the building models and the control logic developed in Tridium to provide load shedding from the simulated buildings.

4.1.2 Building Model Development

We used three U.S. Department of Energy (DOE) reference building models (Deru et al. 2011) indicated in Figure 5 for the lab test. The buildings are modeled in EnergyPlus®,[1] DOE's open-source whole-building energy modeling engine.

[1] https://energyplus.net/.

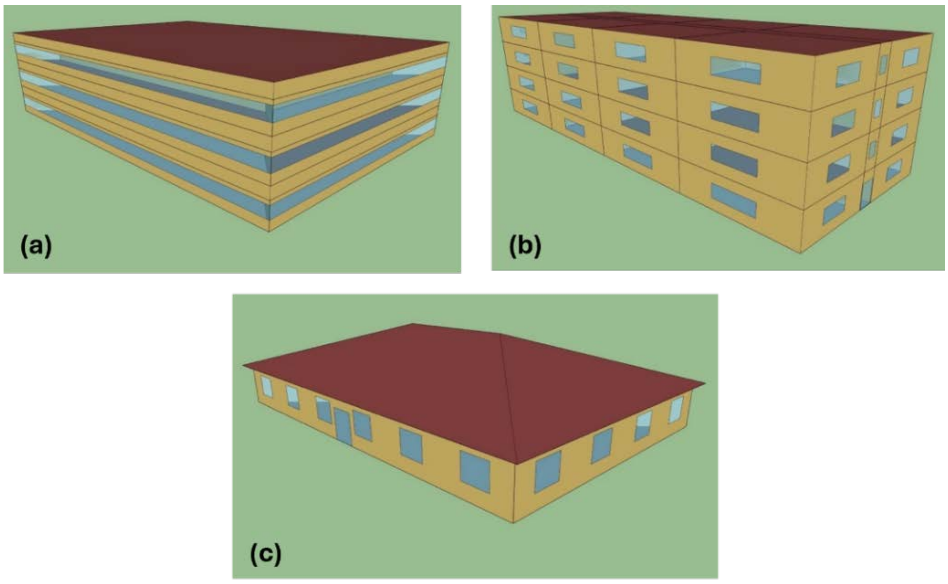

Figure 5. DOE reference building models: (a) medium office; (b) midrise apartment; (c) small office

4.1.2.1 Measure Development

The DOE reference building models are provided as EnergyPlus input data files. However, input data files will not run in Alfalfa by default. They must be encapsulated in an OpenStudio workflow. Here, we use OpenStudio Measures[2] executed in a specific order (e.g., an OpenStudio workflow) to take the energy models in input data file format and translate them for use in Alfalfa. The first measure in the workflow adds desired Alfalfa inputs and outputs, that is, model parameters that we want to actuate with Alfalfa (inputs) and model parameters that we want to monitor (outputs). The next measure takes the DOE reference building model input data files and translates them into an OpenStudio workspace model that Alfalfa can ingest. The final measures in the workflow add the EnergyPlus energy management system code to allow Alfalfa inputs to override the behavior of building models. EnergyPlus interprets and executes the EnergyPlus energy management system code program as the model simulation is executed to provide simulation-time feedback for the Alfalfa inputs.

4.1.2.2 Model Features and Shedding Potentials

To understand the load characteristics of each building, annual simulation data produced by EnergyPlus were analyzed. The peak power consumption and the time of peak were determined for each building for each month, as shown in Table 2.

[2] https://nrel.github.io/OpenStudio-user-documentation/getting_started/about_measures/.

Table 2. Monthly Peak Electric Power Demand

	Small Office		Medium Office		Midrise Apartment		Aggregate Buildings	
Month	Peak kW	Time of Peak	Peak kW	Time of Peak	Peak kW	Time of Peak	Peak kW	Time of Peak
January	15.44	1/12/21 16:53	342.9	1/2/21 6:00	40.7	1/25/21 18:59	380.1	1/23/21 6:00
February	13.69	2/7/21 15:27	342.9	2/15/21 6:00	41.0	2/15/21 18:59	381.3	2/16/21 6:00
March	15.17	3/31/21 13:56	342.9	3/3/21 6:00	45.2	3/31/21 18:15	381.4	3/4/21 6:00
April	14.93	4/24/21 15:29	253.6	4/4/21 5:00	45.4	4/29/21 18:43	289.8	4/4/21 5:00
May	17.26	5/23/21 15:59	176.3	5/5/21 5:00	52.2	5/23/21 17:01	228.4	5/23/21 15:59
June	17.82	6/28/21 15:59	181.8	6/29/21 15:29	57.1	6/17/21 17:01	245.2	6/29/21 15:29
July	19.00	7/18/21 15:59	217.0	7/18/21 15:22	60.4	7/18/21 17:01	287.3	7/18/21 15:25
August	18.53	8/1/21 15:49	214.7	8/1/21 15:30	57.7	8/1/21 17:01	282.5	8/1/21 15:30
September	17.48	9/7/21 15:01	238.7	9/13/21 5:00	53.7	9/7/21 18:17	274.3	9/13/21 5:00
October	16.40	10/13/21 15:05	293.1	10/30/21 5:00	49.7	10/13/21 17:21	329.7	10/30/21 5:00
November	15.58	11/6/21 16:50	342.9	11/27/21 6:00	43.0	11/2/21 17:01	379.6	11/27/21 6:00
December	15.86	12/6/21 16:32	342.9	12/18/21 6:00	40.8	12/28/21 18:59	380.8	12/29/21 6:00
Max	**19.00**	**7/18/21 15:59**	**342.9**	**12/18/21 6:00**	**60.4**	**7/18/21 17:01**	**381.4**	**3/4/21 6:00**

Figure 6 shows the hourly energy use profile by end use for the medium office building on a summer day. For the lab demonstration, we focused on using lighting and interior equipment to provide load shedding; Table 3 summarizes the maximum load shedding potential from these devices for each building type. Both the lighting and interior equipment electricity consumption are operated on a static schedule, in which each building is operating on a set percentage of the maximum electricity consumption, depending on the hour of the day (0–24) and the day of week. These consumption percentage set points can be changed at any time, allowing us to shed a flexible amount of the current consumption.

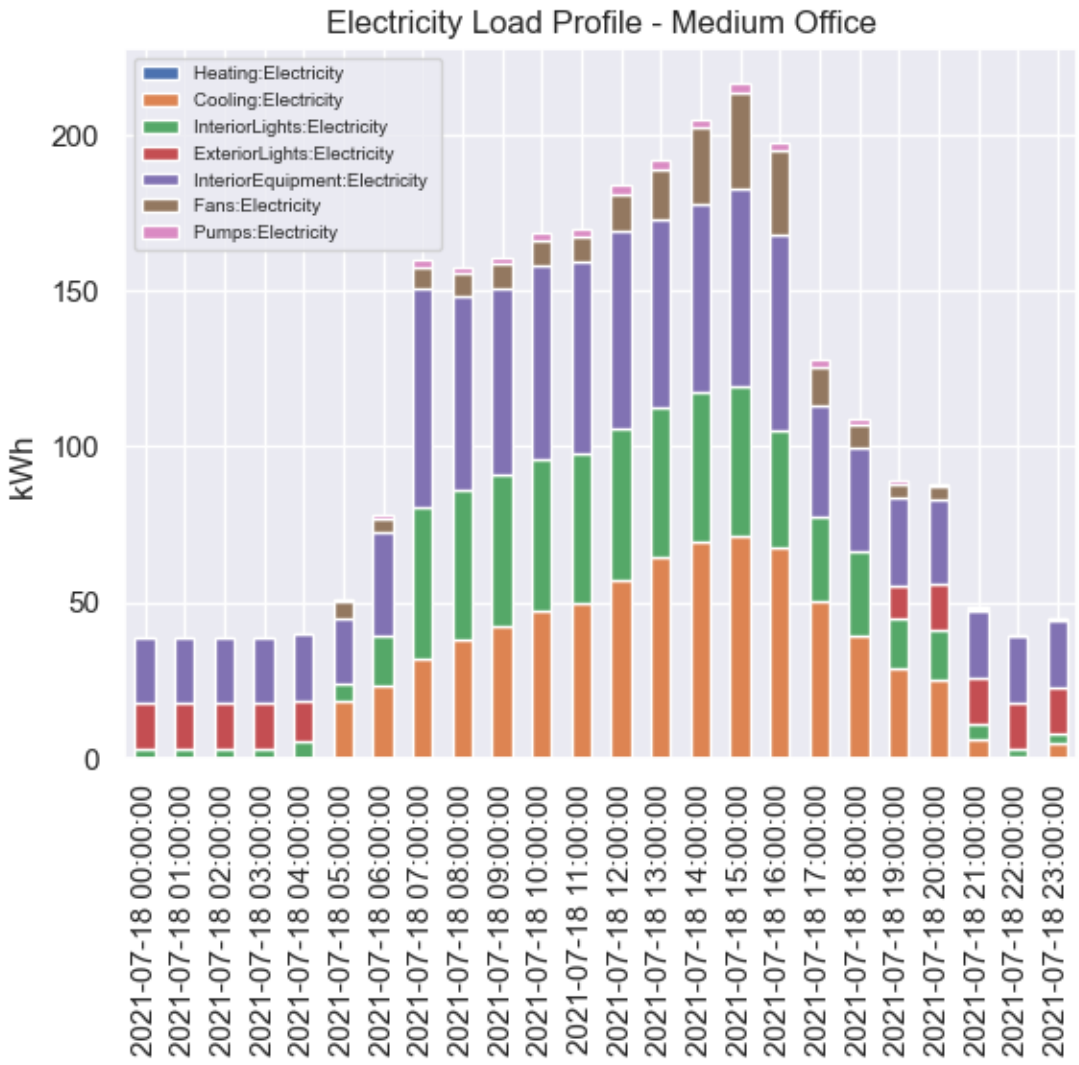

Figure 6. Hourly variation of end-use electricity demand for medium office building

Table 3. Maximum Load Shedding Potential

Building	Lighting Maximum Power (kW)	Interior Equipment Maximum Power (kW)
Small office	5.5	5.5
Medium-size office	53.6	53.6
Midrise apartment	5.1	5.1

4.1.3 Identification of Control Points

Alfalfa exposes each of the virtual buildings as a BACnet device to the Tridium BMS, as shown in Figure 7. Figure 8 shows sample data points from the medium-size office building simulation model that are exposed as BACnet points.

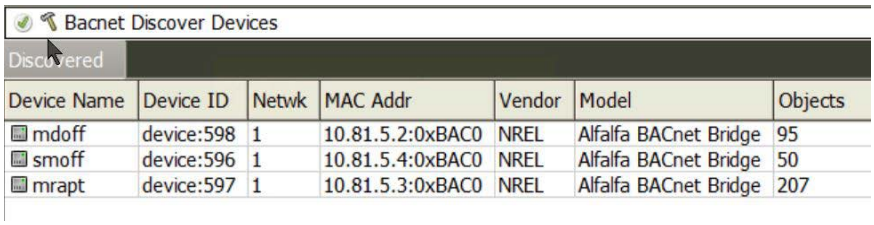

Bacnet Discover Devices

Discovered

Device Name	Device ID	Netwk	MAC Addr	Vendor	Model	Objects
mdoff	device:598	1	10.81.5.2:0xBAC0	NREL	Alfalfa BACnet Bridge	95
smoff	device:596	1	10.81.5.4:0xBAC0	NREL	Alfalfa BACnet Bridge	50
mrapt	device:597	1	10.81.5.3:0xBAC0	NREL	Alfalfa BACnet Bridge	207

Figure 7. BACnet device discovery in Tridium

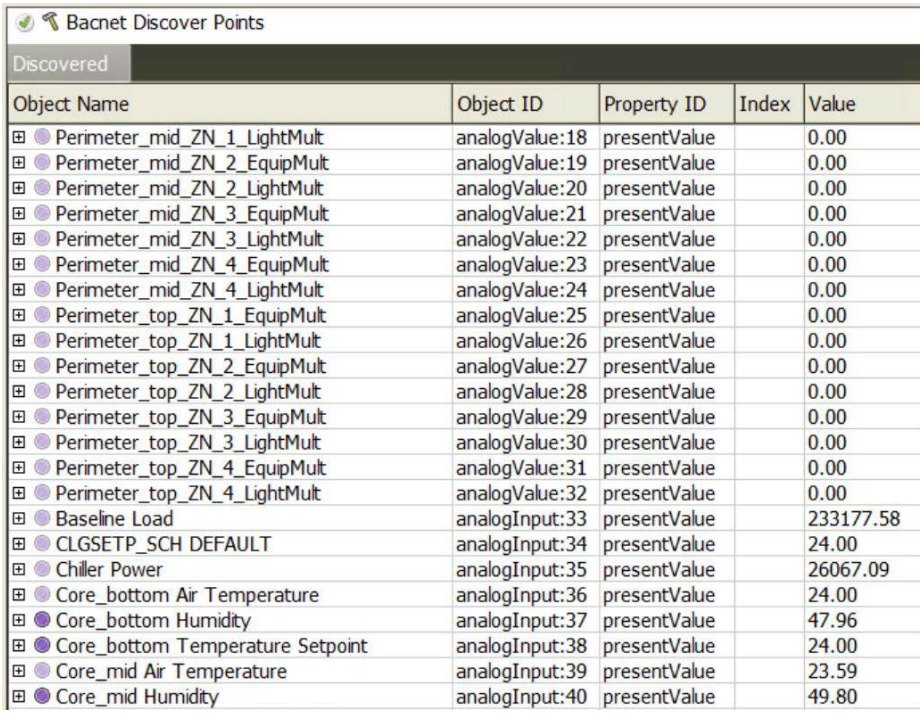

Bacnet Discover Points

Discovered

Object Name	Object ID	Property ID	Index	Value
Perimeter_mid_ZN_1_LightMult	analogValue:18	presentValue		0.00
Perimeter_mid_ZN_2_EquipMult	analogValue:19	presentValue		0.00
Perimeter_mid_ZN_2_LightMult	analogValue:20	presentValue		0.00
Perimeter_mid_ZN_3_EquipMult	analogValue:21	presentValue		0.00
Perimeter_mid_ZN_3_LightMult	analogValue:22	presentValue		0.00
Perimeter_mid_ZN_4_EquipMult	analogValue:23	presentValue		0.00
Perimeter_mid_ZN_4_LightMult	analogValue:24	presentValue		0.00
Perimeter_top_ZN_1_EquipMult	analogValue:25	presentValue		0.00
Perimeter_top_ZN_1_LightMult	analogValue:26	presentValue		0.00
Perimeter_top_ZN_2_EquipMult	analogValue:27	presentValue		0.00
Perimeter_top_ZN_2_LightMult	analogValue:28	presentValue		0.00
Perimeter_top_ZN_3_EquipMult	analogValue:29	presentValue		0.00
Perimeter_top_ZN_3_LightMult	analogValue:30	presentValue		0.00
Perimeter_top_ZN_4_EquipMult	analogValue:31	presentValue		0.00
Perimeter_top_ZN_4_LightMult	analogValue:32	presentValue		0.00
Baseline Load	analogInput:33	presentValue		233177.58
CLGSETP_SCH DEFAULT	analogInput:34	presentValue		24.00
Chiller Power	analogInput:35	presentValue		26067.09
Core_bottom Air Temperature	analogInput:36	presentValue		24.00
Core_bottom Humidity	analogInput:37	presentValue		47.96
Core_bottom Temperature Setpoint	analogInput:38	presentValue		24.00
Core_mid Air Temperature	analogInput:39	presentValue		23.59
Core_mid Humidity	analogInput:40	presentValue		49.80

Figure 8. Medium office building BACnet points discovery in Tridium

Not all the BACnet points are needed for the experiment. The BACnet control points that are used from each building are all the different zones: Light Multipliers, Equipment Plug Load Multipliers, Baseline Load, Whole Building Electricity, Lighting Schedule, and Plug Load Schedule.

4.1.4 Load Shedding Control

The control logic for load shedding is developed in Tridium. As discussed in the previous section, we only considered lighting and interior equipment (plug loads) as the two controllable loads for shedding. As discussed in Section 4.1.2.2, the electricity used by lighting and plug loads varies throughout the day based on preset schedules in the model. These preset schedules are analogues to light dimming schedules in actual buildings. The available load that can be shed at time *t* from a building is calculated as:

$$TSL_t = SL_{Lighting,t} + SL_{Int.equipment,t} \quad (1)$$

$$\boldsymbol{SL}_{\boldsymbol{Lighting,t}} = \sum_{\boldsymbol{i=1}}^{\boldsymbol{N}} \boldsymbol{LS_{i,t}\ X\ LI_i\ X\ SA_i}$$

$$\boldsymbol{SL}_{\boldsymbol{Plug\ loads,t}} = \sum_{\boldsymbol{i=1}}^{\boldsymbol{N}} \boldsymbol{IES_{i,t}\ X\ IESI_i\ X\ SA_i}$$

where *TSL* is total sheddable load, *DL* is sheddable load (kW), *LS* is lighting schedule, *LI* is lighting intensity (in kilowatts per square feet [kW/ft^2]), *IES* is interior equipment schedule, *IEI* is interior equipment intensity (kW/ft^2), *SA* is serving area (ft^2), *t* is time, and *N* is the total number of zones in the buildings.

The preset lighting and plug load schedules from the building simulation model are communicated to the Tridium controller via Alfalfa as BACnet points. The shed request from the MGC is determined as the difference between the allowable load and the building load as indicated in Equation (2).

$$\boldsymbol{SR_t = BL_t - AL_t} \quad (2)$$

In Equation (2), *SR* is shedding request (kW), *BL* is the building load (kW) from the building model that is communicated to Tridium as BACnet point, and *AL* is the allowable load (kW) issued by the MGC.

To provide the required load shedding request, a multiplier (Equation (3]) with a value between 0 and 1 is added to the lighting and plug load schedules.

$$\boldsymbol{multipier_t} = \frac{\boldsymbol{(TSL_t - SR_t)}}{\boldsymbol{TSL_t}} \quad (3)$$

For simplicity, a single multiplier is added for both schedules. A multiplier of 0 is equivalent to turning off the entire lighting and plug loads, whereas a value of 1 means working based on the preset schedules. If the shed request is greater than the total sheddable load, it means the available sheddable load is not enough to provide the requested shedding, and the MGC should operate in a conventional way by opening some of the remotely controlled switchgears to safely run the microgrid. The actual power demand by lighting and plug loads after application of a multiplier is calculated as indicated in Equation (4).

$$APD_{Lighting,t} = SL_{Lighting,t}\ X\ multipier_t \quad (4)$$

$$APD_{Int.Equipment,t} = SL_{\ Int.Equipement,t}\ X\ multipier_t$$

In Equation (4), *APD* is actual power demand (kW).

The multiplier is a numeric writeable BACnet point in the Tridium controller, and its value is communicated to the building model at each time step through Alfalfa. The default value for the multiplier is 1, meaning the system will operate the way it normally operates without load shedding request. Note a single multiplier is used for the lighting and plug loads, but the logic could be easily modified to account for different multipliers.

After application of a multiplier lower than 1, and if the allowable load is higher than the building load plus the hysteresis value defined by the MGC, the multiplier will be recovered back to 1. However, if the allowable load is greater than the building load, but the allowable load is lower than the building load plus the hysteresis, the multiplier value will keep its previous time step value.

Figure 9 illustrates the control logic implemented in the Tridium controller.

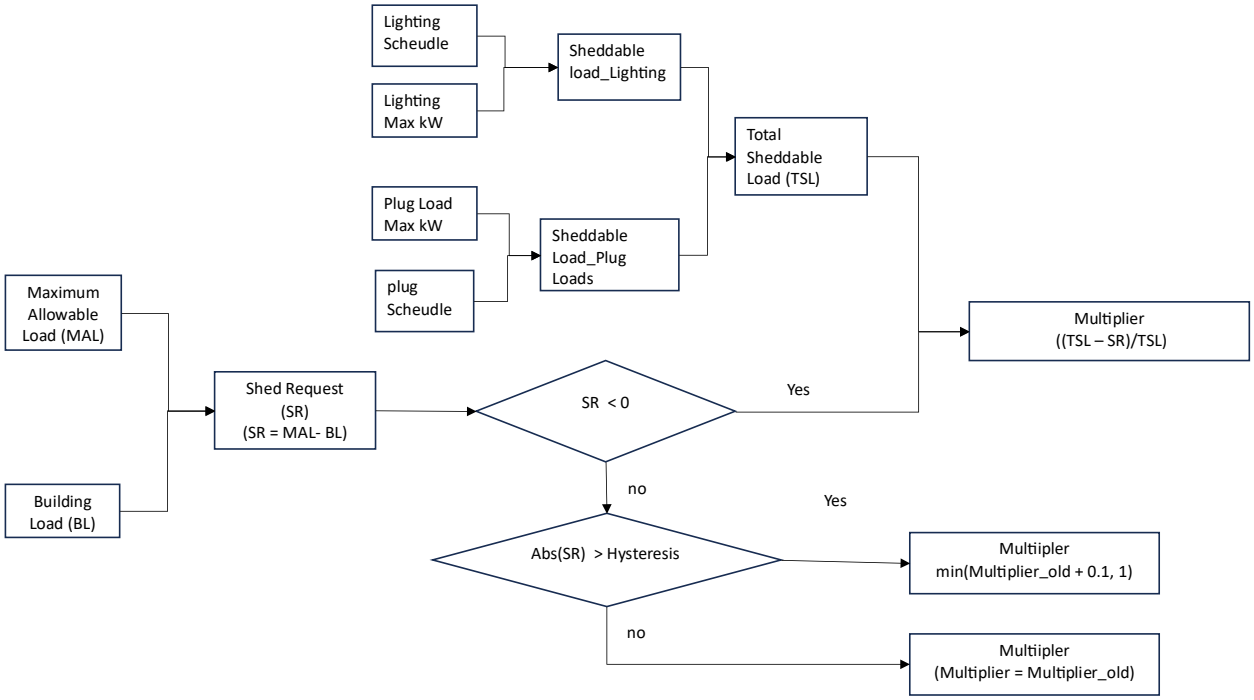

Figure 9. Control logic for load shedding

4.1.5 Laboratory Test Results

The functionality and cybersecurity tests were conducted at ESIF at NREL. The lab setup is similar to that depicted in Figure 4, and the following sections detail the results.

A reference architecture was developed to test the custom equations and processes in a lab setup. This architecture is shown in Figure 10, where there are three loads, one of which (OAD-Load-01) is the bulk load controlled through OpenADR, and the other two are loads controlled in a more traditional fashion by operating the corresponding isolation switch. The power flow through the switches in Figure 10 vary according to the building loads as shown in Figure 11.

OAD-Load-01 is the aggregated load from the three building models (medium-size office, small office, and midrise appartement) discussed in Section 4.1.2, whereas Load 2 is considered a critical load—that is, nonsheddable—and is based on the large office load shape. Load 3 is a modified version of the office load scaled down to provide a noncritical sheddable load.

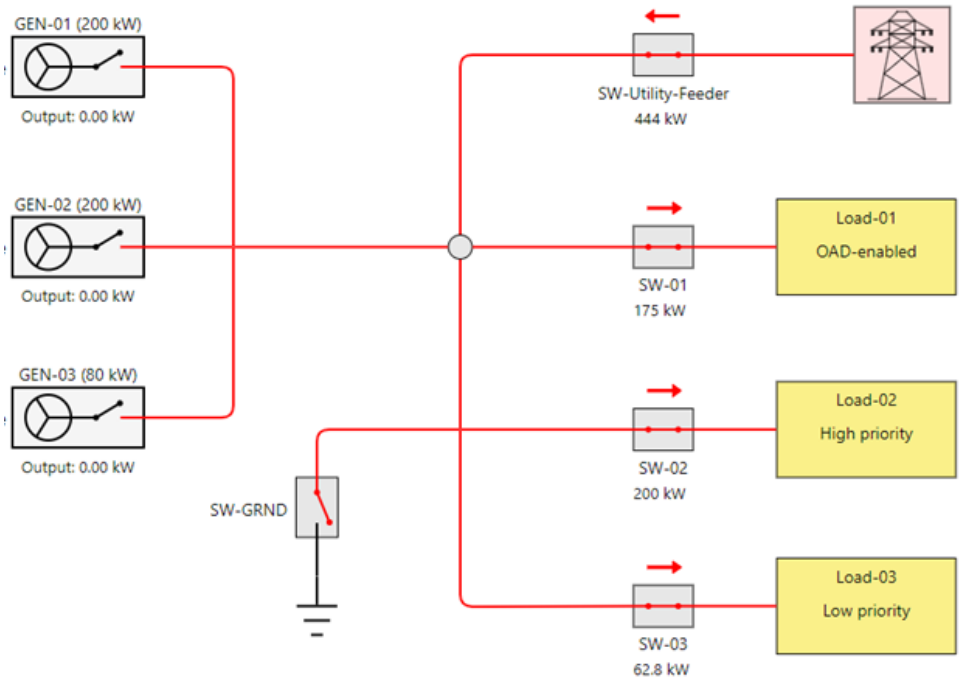

Figure 10. Grid topology for testing and validation

As indicated in Figure 11, the aggregated peak load of this grid is a little less than 390 kW. Three generators were added to this grid to meet this load. Two of the generators are rated at 200 kW each, and one is rated at 80 kW, for a total generation capacity of 480 kW. The microgrid should be able to survive the loss of the 80-kW generator without issue, but loss of either of the 200-kW generators would drop the total capacity to 280 kW, and the grid would need to shed load between the hours of 07:00 and 16:00. Loss of the 80-kW generator plus one of the 200-kW generators would require load shed between the hours of 07:00 and 17:00. Note these generator capacities were chosen more for their ability to test the algorithms and do not necessarily represent the typical capacities found in a real installation.

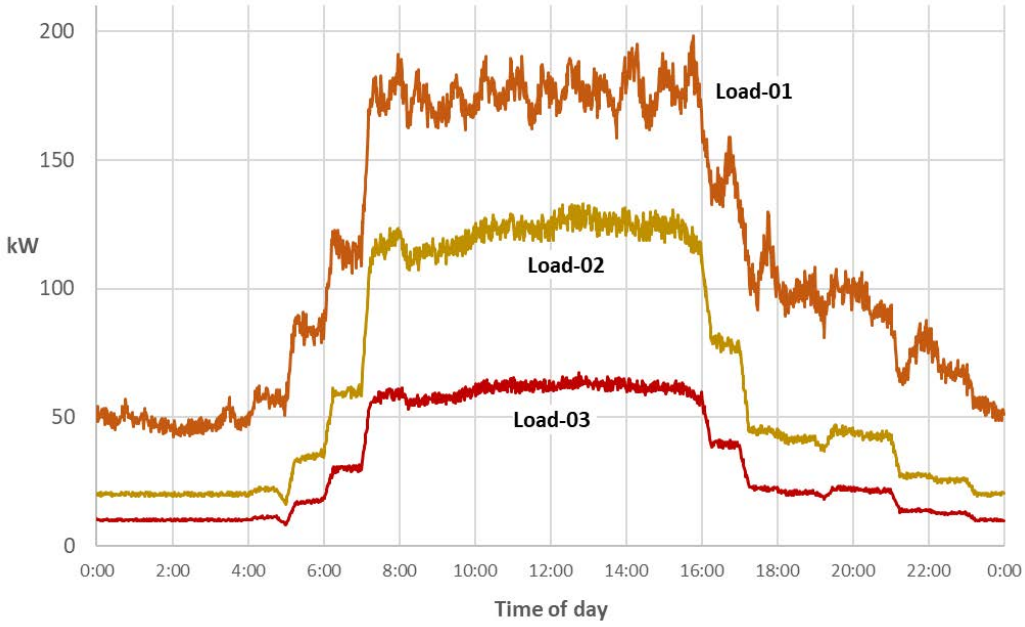

Figure 11. Example summer weekday loads for the reference architecture

During the lab test, we tested six different scenarios that are typical to microgrid operations:

- Grid load approaches grid spinning capacity.
- Large grid source marked unavailable.
- Small grid source marked unavailable.
- Additional capacity available after loads have been shed.
- Transition to grid-tied while loads have been shed by OpenADR.
- Black start with insufficient sources to power entire microgrid.

The results from each scenario are discussed in the following subsections.

4.1.5.1 Lab Test Case 1: Grid Load Approaches Grid Spinning Capacity

This test considers an islanded grid operating at steady conditions. The high-priority load (Load-02) is then increased steadily over a short period, causing the generator output to increase. As the load approaches the available spinning capacity, it becomes necessary to shed load to avoid having the generators trip off on overcurrent. The GridMaster MGC initially requests load reduction from the OpenADR system but sheds lower-priority loads through conventional means if the OpenADR load shed is insufficient.

To demonstrate this case, we used two tests—simulation-only and hardware-in-the-loop. The simulation-only test is a base case simulation test I which the load on Load-02 was increased at a rate of about 50 kW/min, as shown in Figure 12. Once the total load reached around 400 kW, the MGC decreased the allowable Load-01 load value, as shown in Figure 13. As the loads continued to rise, the available load shed dropped to zero. At this point, additional loads had to be shed by conventional methods, and the MGC opened SW-03, causing Load-03 to be de-energized.

After several minutes, the additional load on Load-02 decreased at a rate of approximately 50 kW/min. Because Load-03 was offline, the OpenADR allowable load began to increase at an earlier point and was fully restored once the total load was about 350 kW. As Load-02 continued to decrease, the switch to Load-03 was commanded to close again, and power was restored to that building.

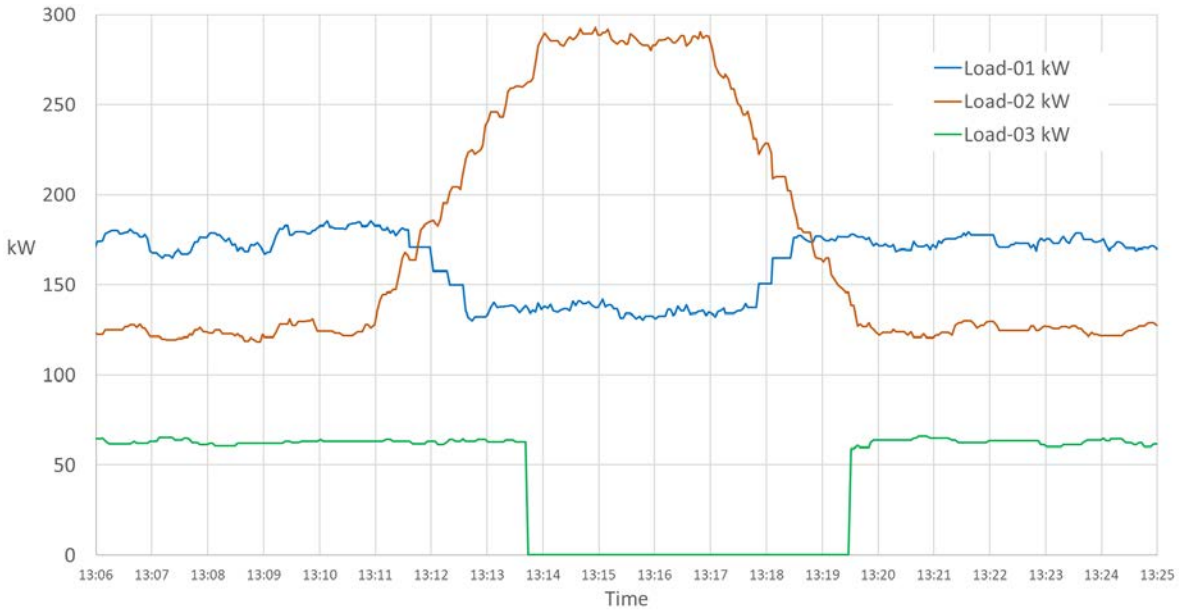

Figure 12. Simulation loads for Test 1 base case

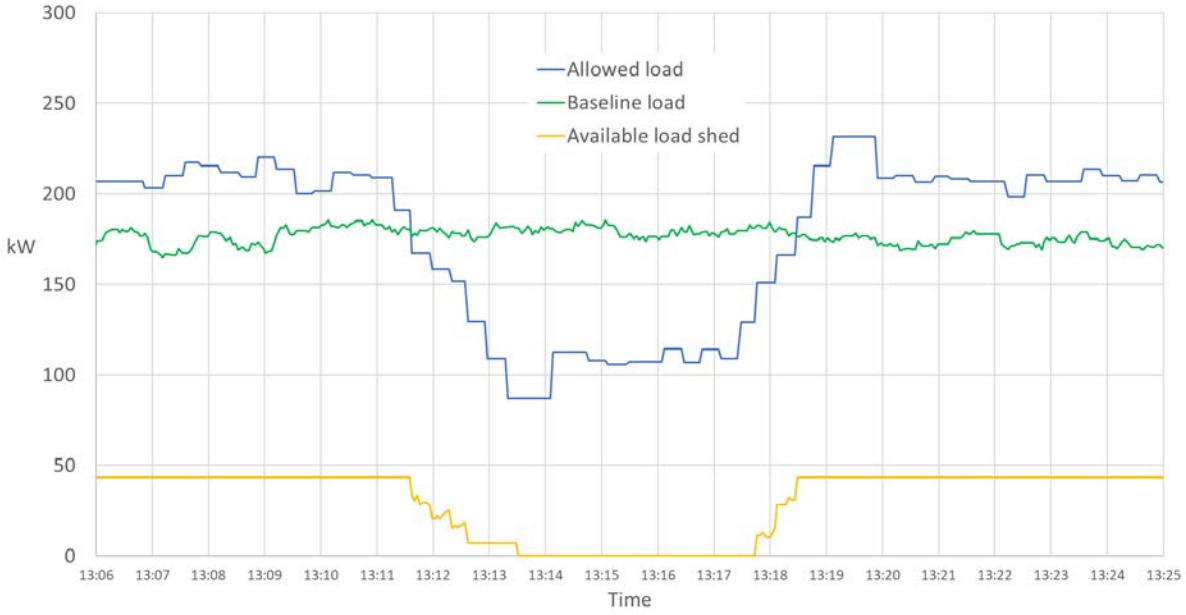

Figure 13. Load management values for Test 1 base case

The hardware in-the-loop test attempted to replicate the simulation-only case results. The setup is indicated in Figure 4, where the building simulation model is manipulated by a control program in a Tridium controller to provide the load shedding request from the MGC. For the simulation-only case, the load shed delay of the native GridMaster algorithms was set to 20 seconds. That is, GridMaster would not attempt to shed loads through conventional means until at least 20 s after the need for load shed was recognized. This delay was to give the OpenADR system enough time to respond to a reduction in the allowable load shed for Load-01.

The results from the initial hardware-in-the-loop test (Test 1) are shown in Figure 14. Load-02 was increased at the rate of approximately 1.0 kW/s, starting at about 120 kW and increasing to about 320 kW. The Load-02 adder remained at this value for 5 min and then decreased at a similar rate until the nominal load was achieved.

In this case, the delay inherent to the simulation setup significantly affected the behavior of the system. The delay was due to two main reasons. The first reason is the 1-min default time interval for polling events from the VTN by VEN. If a load shedding request comes seconds after the last communication, the value will not be communicated to the VEN for the next 1 min. Similarly, if a new report is available at the VEN, the VTN must wait 1 minute to get the value. The other source of delay is the 1-min minimum runtime interval of the building models in EnergyPlus. Once new values are communicated to the building models through Alfalfa, it could be delayed by up to 1 min before it is applied to the models. Alfalfa also needs to wait for another 1 min to read and communicate the values back to the Tridium controller. In total, these delays could add up to a 2-minute delay for communicating the load shedding request from the MGC to the building models and another 2 min for reporting new values back to the MGC. The 2-minute delay in load shedding request means the value received by the building model will be the one in which the load on the building was 120 kW lower than the current value, with the assumed 1-kW/s load increment. As a result, the load shedding provided by the controller will not be enough.

Figure 14 illustrates the issue arising as a result of a delay in the system. As Load-02 increased, the allowed kilowatts for Load-01 began to decrease around 09:52. The controller began adjusting the set point to shed loads (as depicted by a decline in available load shedding in the figure) after a 1-min delay. However, the actual load reduction in the simulation model was not communicated back to the GridMaster for at least 4 min.

Recognizing the loads were escalating beyond acceptable levels, the GridMaster acted by curtailing Load-03 through conventional methods, that is, opening SW-03 and disconnecting the entire building (as shown by the green line in the middle chart). Despite this action, Load-02 continued to rise, and by 09:55, the generators were nearing their rated capacity. The two red ovals on the bottom graph highlight instances when the generators had briefly operated at or above their capacity. Although this is tolerable for short durations, it is generally not considered best practice.

In response, the GridMaster opened SW-02 and shed Load-02 as a measure to alleviate the strain. With this load removed from the system, there was sufficient generation capacity, leading to the restoration of Load-02 after approximately half a minute. This cycle repeated twice more until the OpenADR system finally began managing Load-01 around 09:57. As the power demand for Load-01 started to decrease, the system stabilized, allowing Load-02 to remain energized.

Load-02 began decreasing at 10:00 and continued to decrease for the next 200 s. During this time, GridMaster increased the Load-01 allowed power accordingly, but as before, the simulation model did not respond for several minutes. During this time, GridMaster recognized that there was now sufficient spinning capacity to restore Load-03 and commanded SW-03 to close at approximately 10:01.

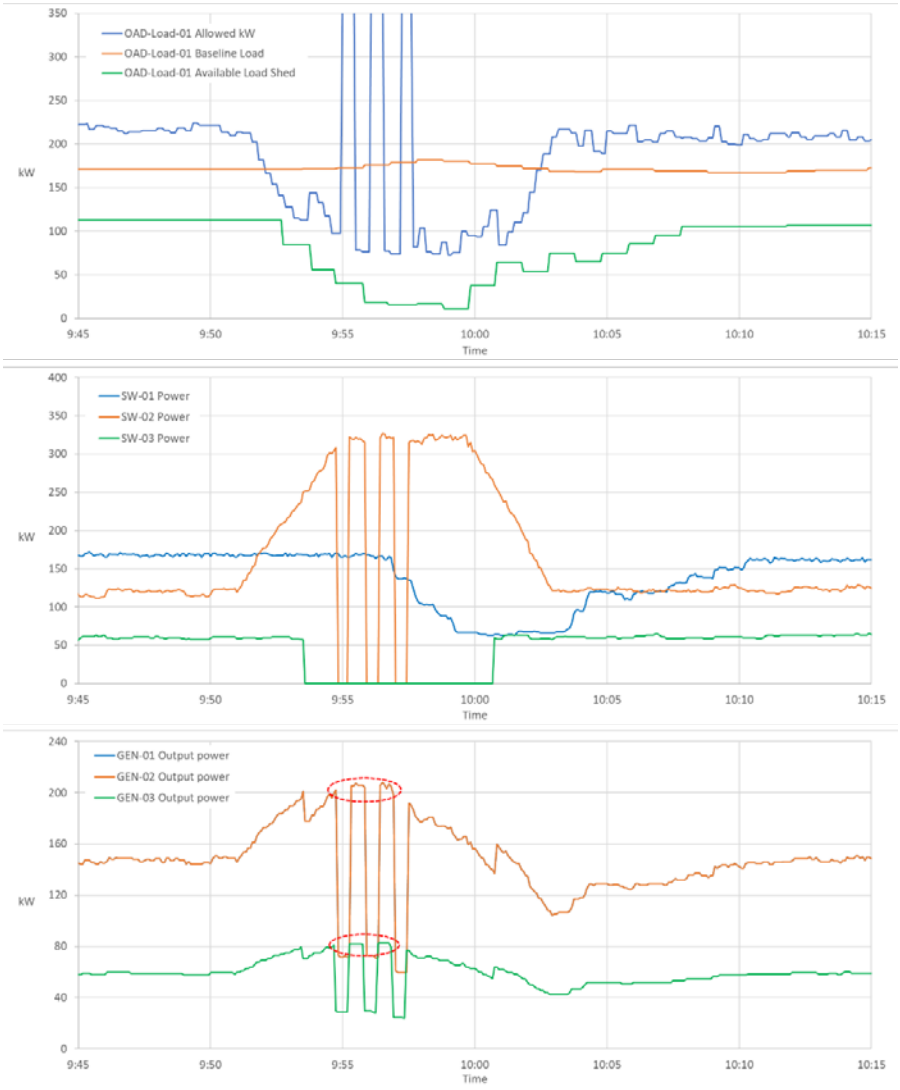

Figure 14. Test 1 results with load shed delay of 20 s and load increase rate of 1.0 kW/s.

Highlighted areas on the bottom graph show where generators operate over capacity; note the overlap of the curves for GEN-01 and GEN-02 output power.

From these results, it is evident that a load ramp rate of 1.0 kW/s is too fast for the inherent OpenADR response time in this simulation. To better gauge the response of the system to the load ramp rate, the test was repeated, but this time using a rate of 0.17 kW/s (i.e., 10 kW every 60 s). The results for this test are shown in Figure 15. In this case, the delay in the simulation setup was manageable because of the slower ramp rate. The generators did not exceed their rated capacity, and Load-02 did not cycle.

To evaluate the impact of the GridMaster load shed delay for conventional curtailment of loads, we used a load shed delay of 90 s. That is, GridMaster would not open switches to curtail loads until 90 s after the need was identified. The goal of this test was to see if the building outages could be minimized or at least delayed, giving the OpenADR controller time to respond. In this case, the longer time delay resulted in Load-03 not curtailing as soon as it did in the previous test. In addition, the system was able to survive a higher load addition without causing the generators to exceed their rated capacity. Load-03 was shed by necessity but only after there was no additional load shed availability on Load-01.

Overall, these tests demonstrate GridMaster can successfully communicate with an OpenADR appliance using the Translator as an intermediary gateway. Although there are challenges associated with the inherent delays in the OpenADR response, these challenges can be addressed by selecting existing tunable parameters within GridMaster. In field applications, the delay caused by the simulation model runtime is not a concern because the actuated values come directly from sensors. Furthermore, the OpenADR 2.0b version permits a custom data polling time interval lower than 1 min, which is the default used in the lab test. Therefore, the observed delays could be significantly reduced in field applications by leveraging this feature.

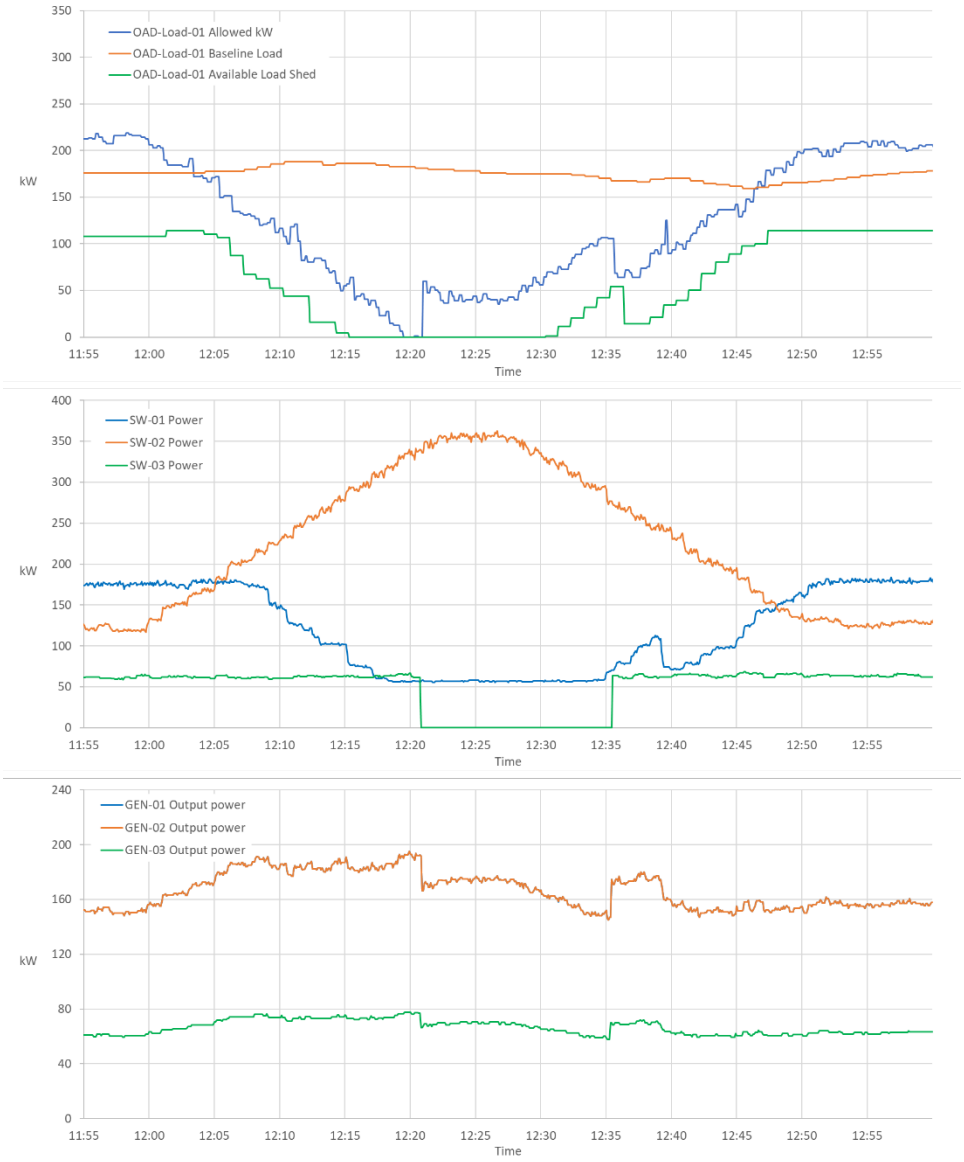

Figure 15. Test 1 results with load shed delay of 90 s and load increase rate of 0.17 kW/s.

Note the overlap of the curves for GEN-01 and GEN-02 output power.

4.1.5.2 Lab Test Case 2: Grid Source Becomes Unavailable

Test Case 2 considers an islanded grid that is operating steadily until one of the generators is suddenly marked unavailable. This may occur if the generator has an internal fault or has to be taken offline for any reason, such as maintenance, fault alarms, or other operational concerns.

Once a generator is marked unavailable, GridMaster attempts to shed enough load to allow the grid to remain energized in the absence of the now-unavailable generator.

This test examine two different scenarios: (1) the small generator is marked as unavailable and (2) the larger generator is marked as unavailable.

A. Large grid source marked unavailable: The first iteration of this test investigated what happens when a large generator is marked as unavailable, and the GridMaster load shed delay was set to 90 s. The timeline for events is shown in Figure 16, where the arrow on the bottom graph shows the point at which GEN-02 was set to unavailable. As soon as the generator availability was set false, the Load-01 allowed the power value to go almost to zero. As before, the actual load was not reduced until almost 5 min later. In the meantime, Load-02 and Load-03 were shed through conventional means less than 2 min after the generator availability was set to false. GridMaster then turned the generator off about 15 s after the loads were shed.

This test was repeated with the GridMaster load shed delay reduced to 20 s. The results, shown in Figure 17, are similar to those of the previous test, with the exception that Load-02 and Load-03 were dropped about a minute earlier (as would be expected given the reduction in the load shed delay). This test has an interesting manifestation of the delayed OpenADR response: even though the allowed power for Load-01 was initially reduced, it was then restored to its original value because Load-02 and Load-03 were dropped before Load-01 was acted upon by OpenADR. The effect was that Load-01 never experienced a load shed event.

Figure 16. Test 2A results showing loss of availability of large generator and load shed delay of 90 s

Figure 17. Test 2A results showing loss of availability of large generator and load shed delay of 20 s

B. Small grid source marked unavailable: The tests described in Test Case 2A were repeated, this time gauging the effects of having the small generator marked as unavailable. We expected the results would be similar, except that only Load-03 would have to be shed because a smaller source was being removed from service.

As before, the first test had a GridMaster load shed delay of 90 s. The resulting data, shown in Figure 18, show the ***Load-01 allowed kW*** was reduced within a few seconds after the generator was set to unavailable. As before, Load-01 does not start to decrease until more than 3 min after the allowed power value is lowered. At this point, there was an interesting dynamic between GridMaster and OpenADR systems. Eventually GridMaster curtailed Load-03 by opening switch SW-03, at which point the Load-01 allowed power was restored to its original value. But in the meantime, the original reduction of the allowed load was acted upon, and Load-01 began to decrease. Eventually Load-01 was reduced to the point where there was sufficient spinning capacity to allow Load-03 to be restored. This is precisely what happened (shortly after 13:35), and as soon as it did, the ***Load-01 allowed kW*** value was once again reduced. At this point, Load-03 remained energized while Load-01 remained in a partial load shed condition.

This test was repeated but with the GridMaster load shed delay set to 20 s. In this case, the dynamic between GridMaster and the OpenADR simulation is different. Like the previous test (and as shown in Figure 19), as soon as the generator was set to unavailable, the Load-01 allowed power was reduced. But this time, Load-03 was shed within half a minute of the generator being set to unavailable, and the allowed power was restored to its nominal value within a minute. The delay in response to the OpenADR means that Load-01 was never reduced, and therefore the overall load on the generators never dropped to the point where Load-03 could be restored.

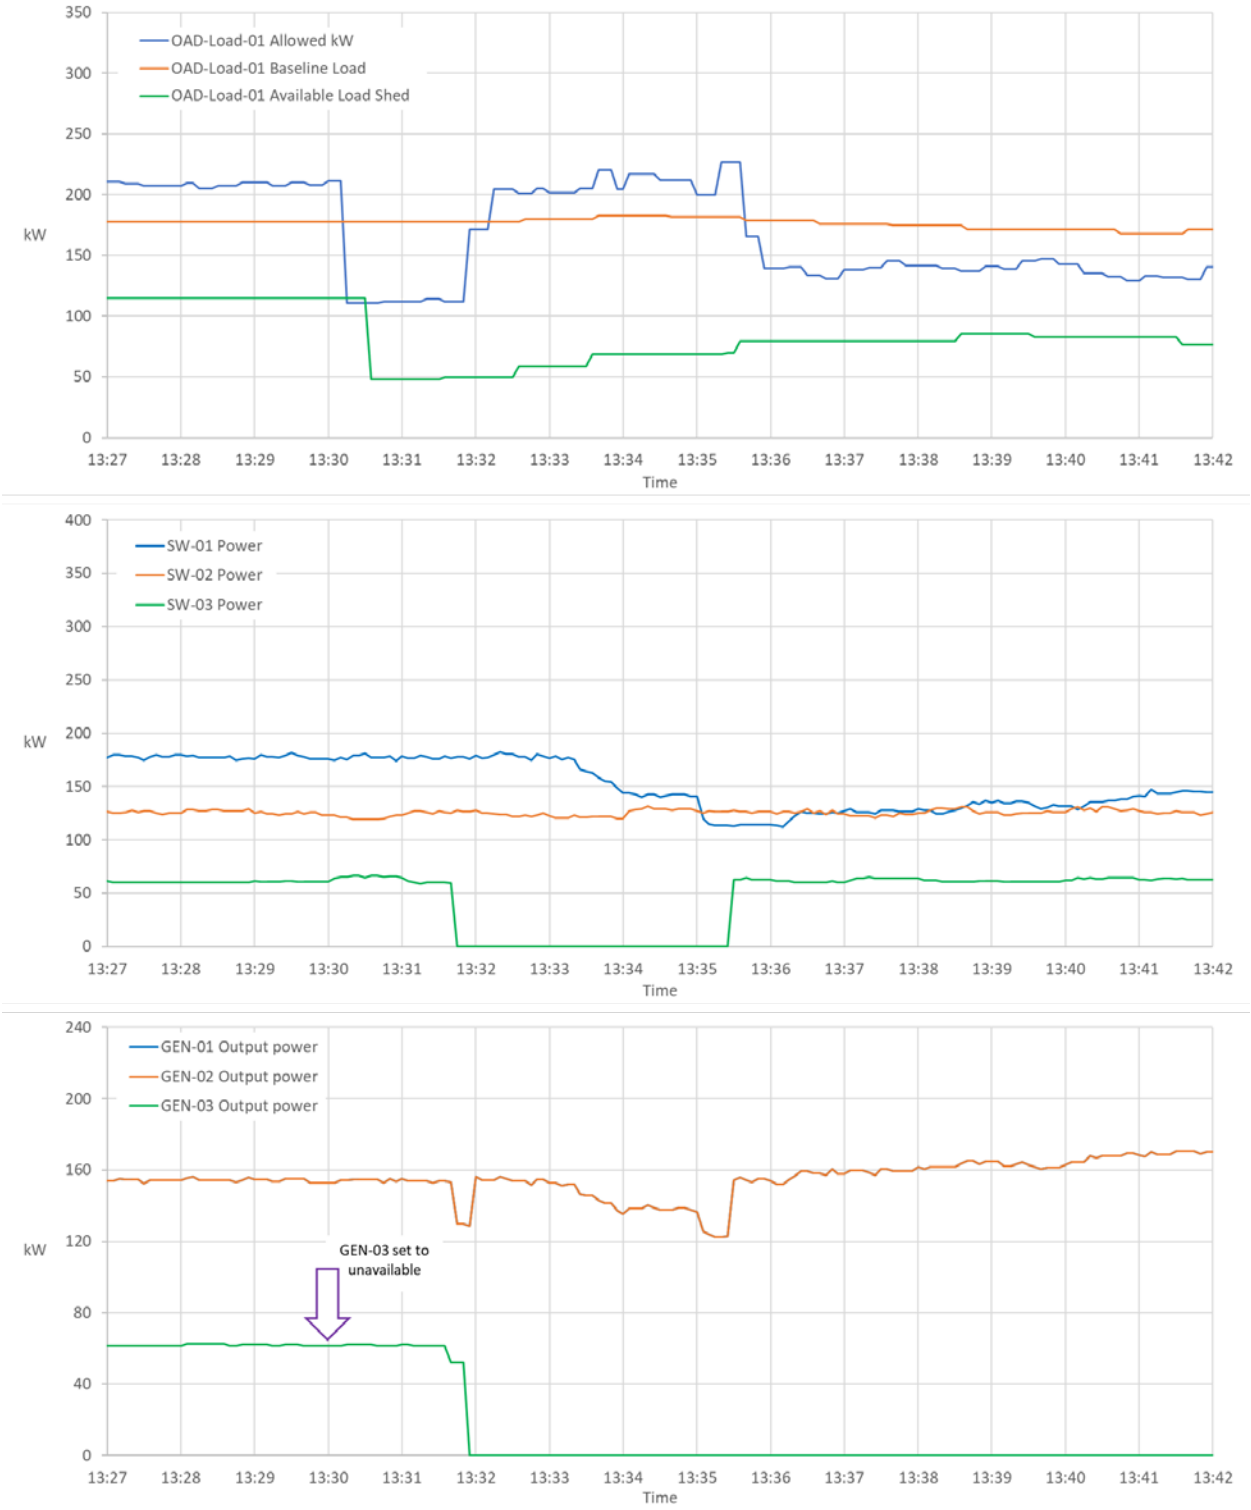

Figure 18. Test 2B results showing loss of availability of small generator and load shed delay of 90 s.

Note overlap of curves for GEN-01 and GEN-02 output power.

Figure 19. Test 2B results showing loss of availability of small generator and load shed delay of 20 s.

Note overlap of curves for GEN-01 and GEN-02 output power.

4.1.5.3 Lab Test Case 3: Additional Capacity Available After Loads Have Been Shed

These tests examine the ability of the system to restore loads once they have been shed but then additional generation capacity is made available. In this case, the GridMaster load shed delay period should not have any impact on the results. The only delay seen in the restoration of loads will be the inherent latency in the OpenADR portion of the simulation.

This test was performed twice. The first time was with one of the larger generators marked as unavailable. This caused Load-02 and Load-03 to be shed, although Load-03 was eventually restored because it has a lower anticipated load. This starting condition is shown in Figure 20, where Load-01 is subject to a load reduction of about 25 to 30 kW. Once GEN-02 was marked as available, it turned on within about half a minute, and Load-02 was restored within about 75 s (see Figure 21). The OpenADR controller restored Load-01 over the following 6 min.

The second iteration of this test started with marking the small generator unavailable. All three loads remained energized, but the Load-01 allowed power was set to about 30 to 40 kW below the baseline load. Figure 22 shows what happened after the generator was made available again. As before, the generator turned on within half a minute, the Load-01 allowed power was set above the baseline load, and the OpenADR controller restored the load gradually over the next 5 to 6 min.

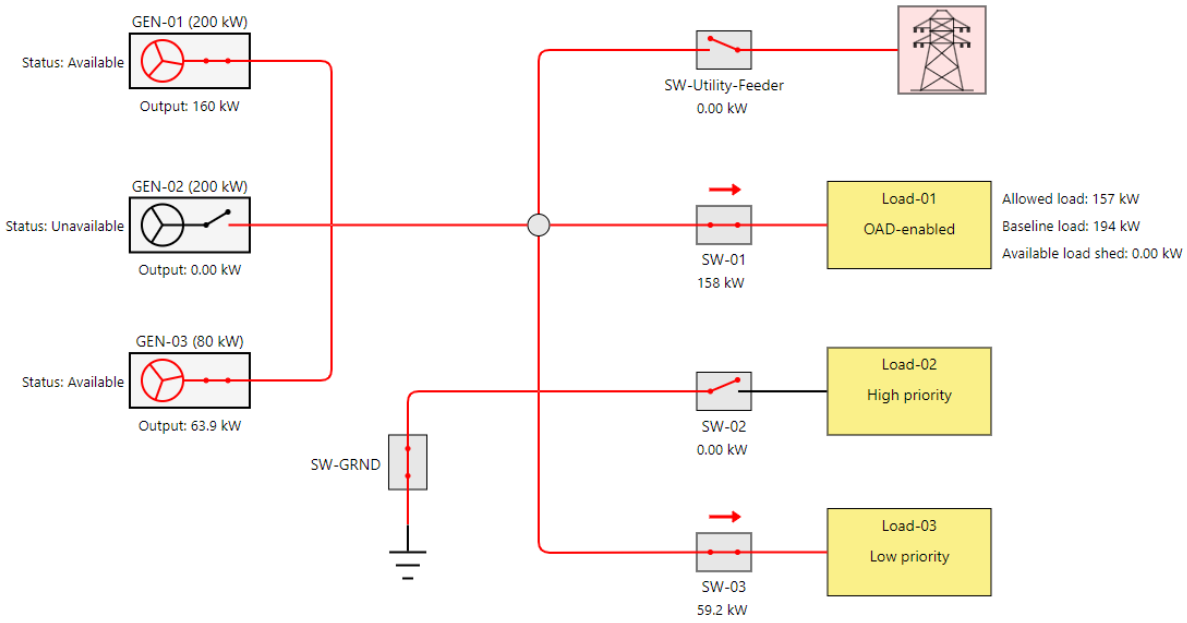

Figure 20. Initial conditions for the capacity addition test

Figure 21. Test 3 results showing addition of large generator capacity

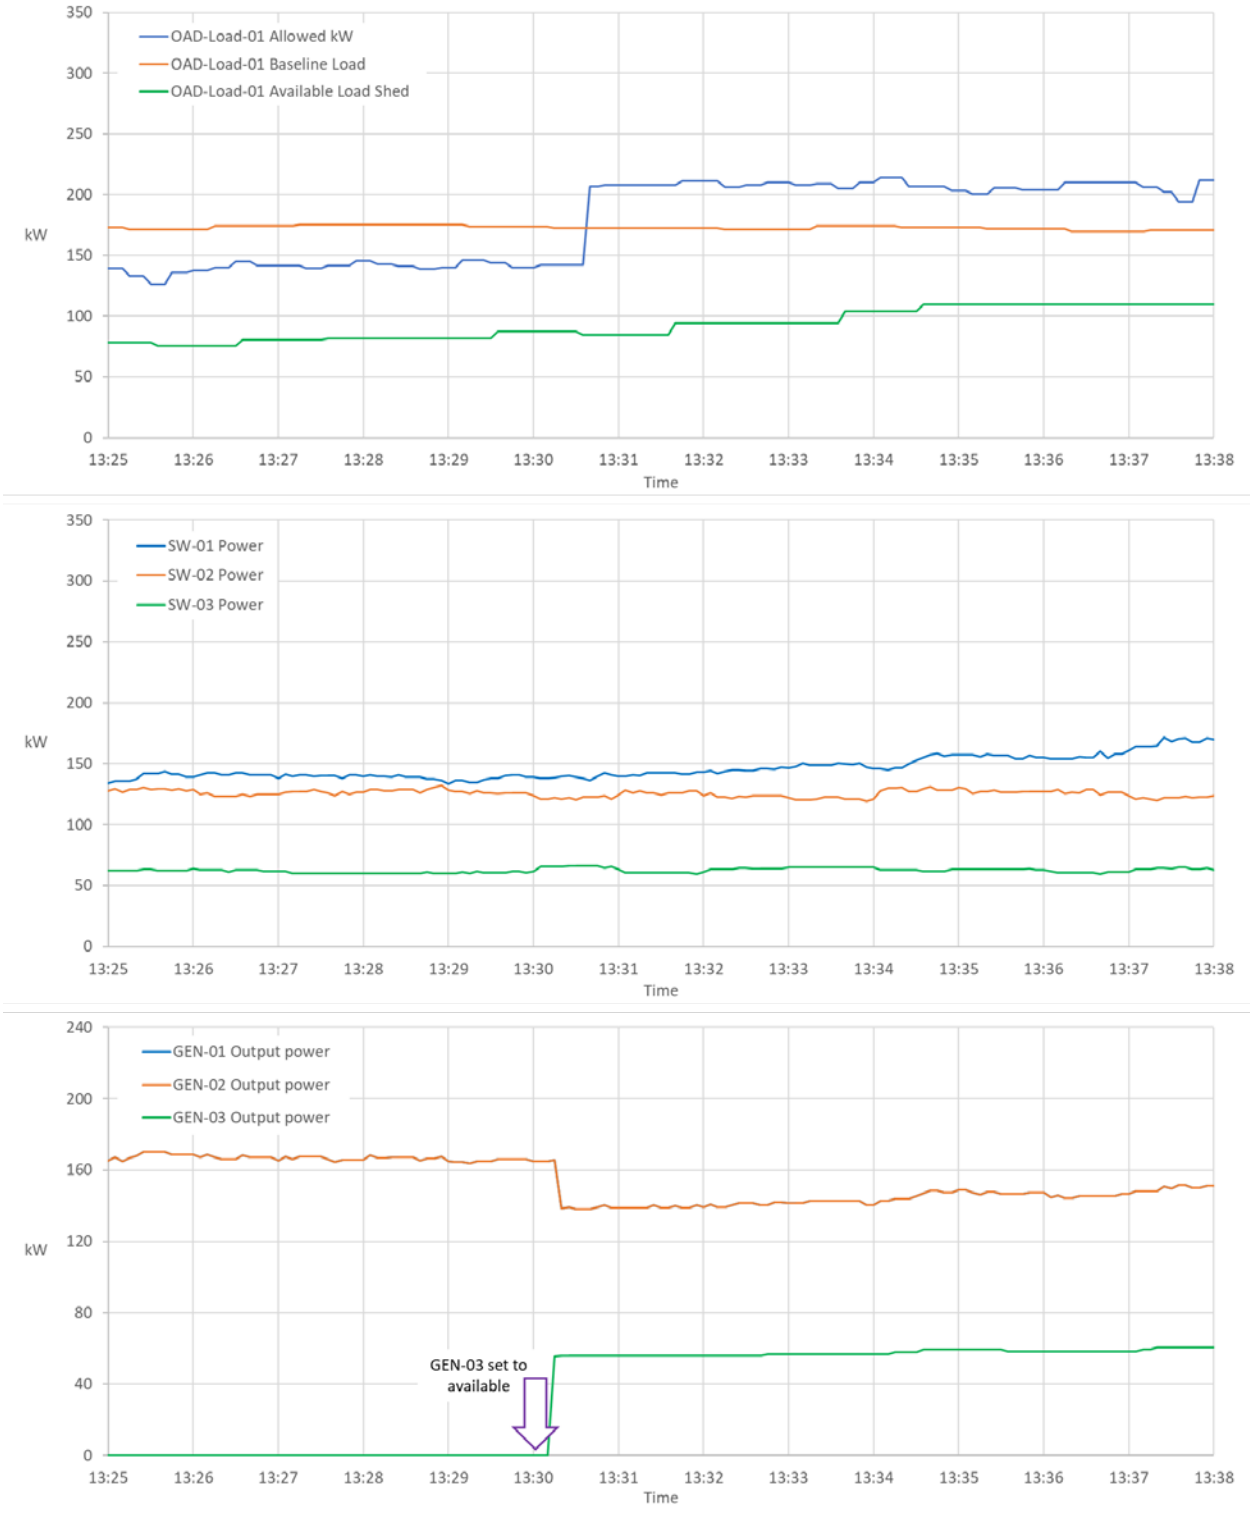

Figure 22. Test 3 results showing addition of small generator capacity.

Note overlap of curves for GEN-01 and GEN-02 output power.

4.1.5.4 Lab Test Case 4: Transition to Grid-Tied While Loads Have Been Shed by OpenADR

Test Case 4 investigates the scenario when loads have been curtailed and the microgrid is then reconnected to the utility grid through a closed transition. Once connected to the utility, there is no need to perform load shedding, and it is expected that all loads will be restored.

To partially curtail loads, GEN-03 was set to unavailable, and Load-02 was increased to the point where Load-01 was partially shed through the OpenADR commands and Load-03 was completely shed through the opening of SW-03. Once the system reached a steady state, GridMaster was instructed to perform a closed (i.e., seamless) transition to grid-tied mode. Once this command was sent, the utility switch was closed, the grounding switch was opened, and the generators were turned off. Once the system was grid-tied, the allowed power was set to −1, indicating to the OpenADR controller that it was allowed to restore all loads that had been shed.

The result of the test is shown in Figure 23. The bottom graph in the figure has an additional line (in red) showing the power flow through the utility switch. Once the utility switch is closed, the power flow will increase and continue to do so as the generators turn off. The test result shows that the system behaved as expected. Once the utility switch was closed, all loads were restored either by operating switchgear or through OpenADR. In the former, SW-03 was operated within half a minute of utility connection. In the latter, the OpenADR system took between 5 and 8 min to fully restore Load-01.

Figure 23. Test 4 results showing transition to grid-tied mode

4.1.5.5 Lab Test Case 5: Black Start With Insufficient Sources To Power Entire Microgrid

Test Case 5 considers the case when GridMaster performs a black start, but there may be insufficient generation capacity to handle the anticipated sum of all the building loads. In this case, GridMaster determines if the full load of Load-01 can be added to the microgrid or if the load has to be shed before closing SW-01.

The initial conditions for this test are the microgrid is grid-tied with the utility switch closed and all generators off. GEN-03 is marked as unavailable, and Load-02 is increased to the point OpenADR load shedding will be necessary in an islanded mode.

The test was initiated by requesting GridMaster to perform an open transition to island mode. All generators were initially off. Figure 24 shows the sequence for the black start. Once the utility switch opened, the generators were turned on one at a time and connected to the bus. The loads were then added in prioritized order. Finally, Load-01 was connected and was modified by the OpenADR controller. Typically, Load-03 is then curtailed by operating SW-03 (since GEN-03 was unavailable and did not turn on).

Figure 25 and Figure 26 show the loads in each building and the load management for a simulation test that was used as a base case. In this case, only GEN-01 and GEN-02 were connected since GEN-03 was marked as unavailable. Load-02 was restored first because it had a higher priority than Load-01. Approximately half a minute later, Load-01 was connected. The OpenADR values are shown in Figure 26. Note that once the generators were turned on, the allowed load jumped to a very high value. This is because there were no loads currently connected. Once Load-02 was connected, the allowed load decreased to a value below the baseline load, and Load-01 experienced load shedding once it was connected. This OpenADR load was energized approximately 20 s after the load adjustment was made.

For the actual lab test, as with the previous tests, several iterations were run to evaluate how the OpenADR response affects the results. In the first case, the GridMaster load shed delay was set to 90 s, and an additional 80 kW was added to Load-02. The results are given in Figure 27, which also shows the power flow through the utility switch. In this case, the combination of the long load shed delay and the delayed OpenADR response caused the generators to operate well over their rated capacity for about 2 min. In the real world, it is likely the generators would have experienced an overcurrent fault and shut off. There was also a problem with Load-02 being cycled on and off until the OpenADR load response occurred. Load-03 was initially dispatched but was then curtailed in response to the load on the generators.

To keep the generators from tripping off, the previous test was repeated with a GridMaster load shed delay of 20 s and with 40 kW added to Load-02. The results (Figure 28) show that the lower load prevented Load-02 from cycling, as would be expected, but Load-03 was dispatched only to be curtailed again once the generators exceeded their rated capacity.

In contrast to the base test case (in which the OpenADR response was near-instantaneous), the inherent delay of the OpenADR system in the simulation setup creates an issue with the setting of the demand response prior to a black start. Simply put, the command to have the OpenADR load shed is not implemented fast enough for the system to respond as desired in this scenario.

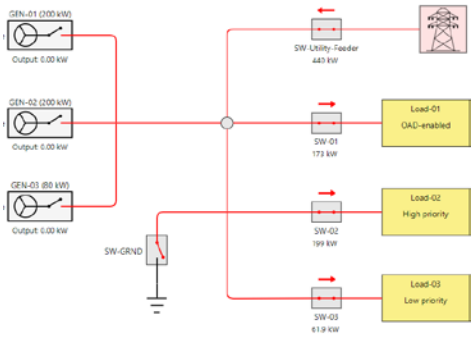

t = 0 s, grid-tied

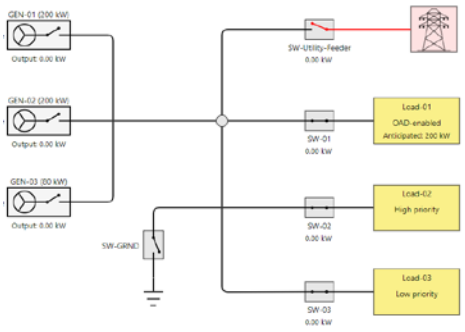

t = 40 s, de-energized

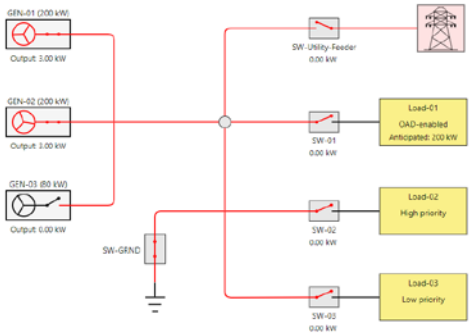

t = 110 s, generators dispatched

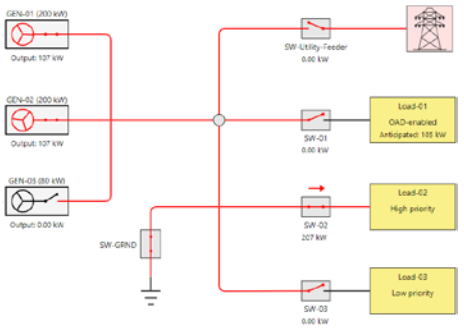

t = 150 s, Load-02 energized

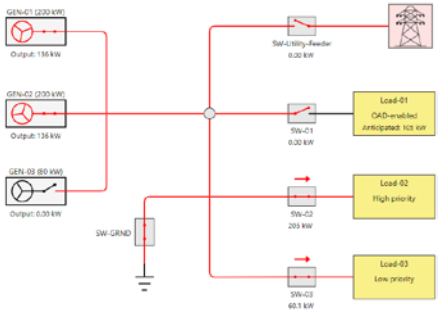

t = 185 s, Load-03 energized

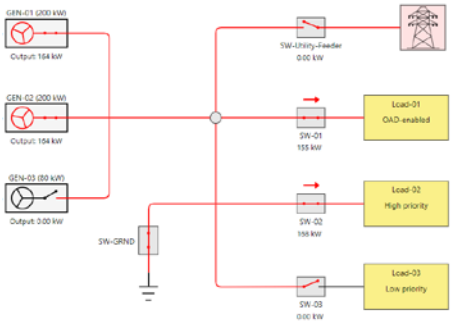

t = 340 s, Load-01 energized

Figure 24. Black start sequence used in Test 5.

Approximate elapsed times since process initiation.

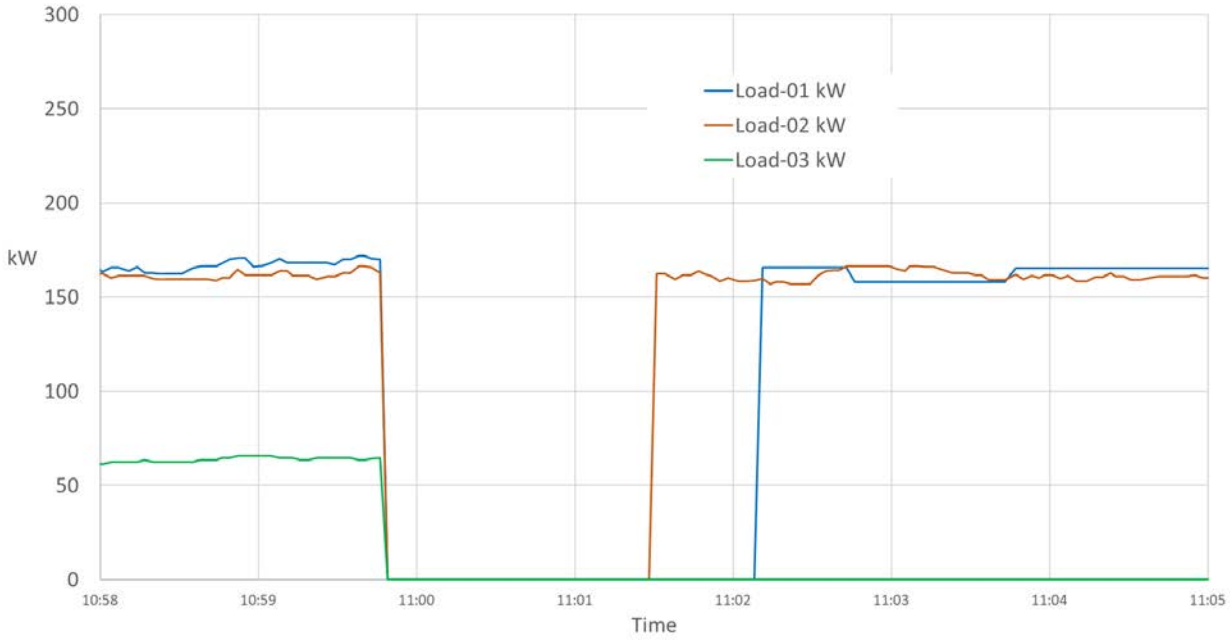

Figure 25. Simulation load for Test 5 base case

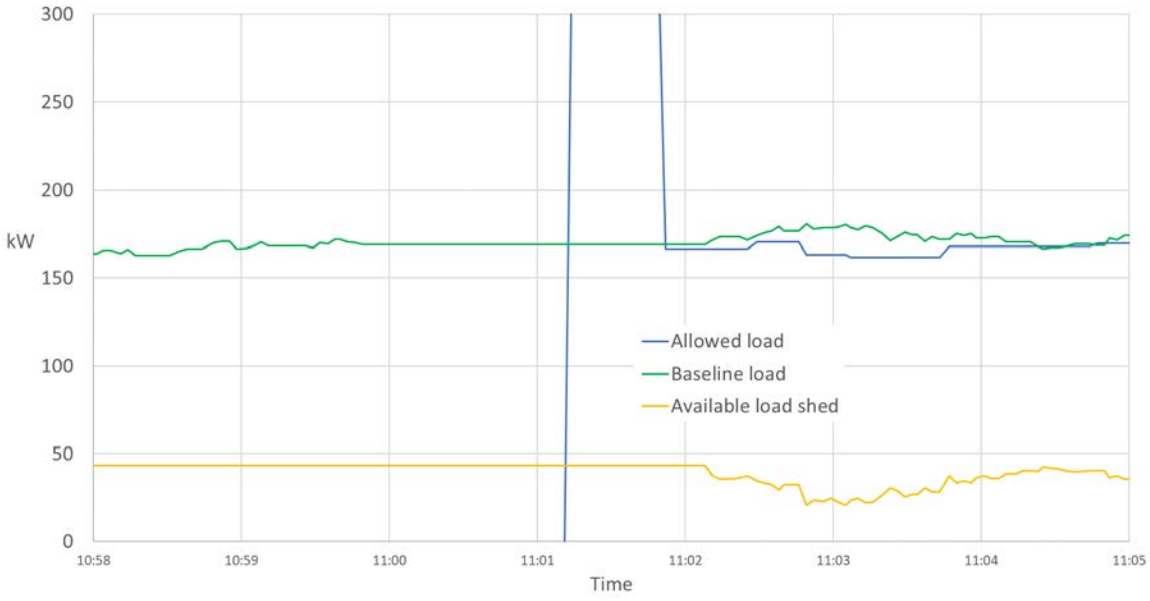

Figure 26. Test 5 base case load management values

Figure 27. Results for Test 5 with load shed delay set to 90 s and a Load-02 adder of 80 kW.

Highlighted areas on bottom graph show where generators operate over capacity; note overlap of curves for GEN-01 and GEN-02 output power.

Figure 28. Results for Test 5 with load shed delay set to 20 s and a Load-02 adder of 40 kW.

Highlighted areas on bottom graph show where generators operate over capacity; note overlap of curves for GEN-01 and GEN-02 output power.

4.1.6 Conclusions From Lab Testing

The lab test investigated the ability of an MGC to communicate with an OpenADR module connected to a Tridium controller. The goal was to see if this was a viable solution for managing loads in buildings that are part of a government microgrid.

Although the tests were largely successful, it must be recognized that there were limitations in the lab setup that were mainly due to delays in the various levels of communication. The expected delays in the various components of the lab setup are indicated in Figure 29. The setup could result in 3- to 5-min delays in the BMS response to the MGC request, causing behaviors indicated in the test results. The field test presented in Section 4.2.4 indicated these delays are significantly reduced during actual field implementation.

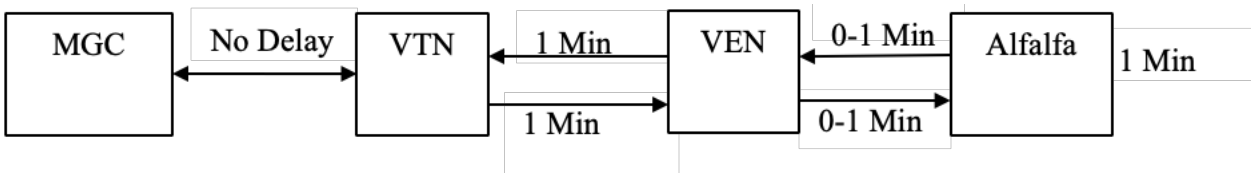

Figure 29. Expected delays in the lab setup

The lab test indicated that the developed methodology could be successfully used for the following operations:

- When islanded, OpenADR could be used to set target power demands for buildings within the microgrid. This would be done at the time of islanding and be based on the mission criticality of each building. A custom process would be created that compares the sum of capacities of all available power sources to the sum of all anticipated building loads. If there is insufficient generation, then OpenADR would be used to limit building loads to manageable values.
- If a system is operating in islanded mode and a generation source has to be removed from service for maintenance, the OpenADR system could be used to reduce overall microgrid load to the point where the generation source can be curtailed. Note, in this case, the loads would have to decrease to the point where the loss of generation source would not cause additional, coarser load shed.

There are several scenarios for which OpenADR is not an appropriate mechanism to implement demand reduction, mostly due to the inherent latency of OpenADR response to a request. These scenarios include:

- Decreasing loads rapidly, when necessary, due to an unexpected loss of a power source.
- Attempting a black start that requires pre-shedding of loads prior to energizing the bus.
- Reducing demand on a microgrid due to a sudden and unexpected increase in load.

4.2 FIELD TEST

This section describes the results from an installation of the prototype system at JB MDL in New Jersey. This was a control hardware-in-the-loop test that mixed both simulated and actual

equipment. A microgrid with simulated generators and switchgear was employed for the testing, along with two simulated loads, one critical and one low-priority load that could be shed through traditional means (i.e., opening a switch that controls power flow to the load). Building 5403 at JB MDL was used as the actual physical load for the tests.

An S&C GridMaster MGC was temporarily installed in the building along with a new JACE 8000 controller. The JACE is designed to support the Niagara controller system and is OpenADR-compliant. The JACE was programmed by NREL to send data and respond to commands from the GridMaster controller.

The chiller plant in Building 5403 was controlled by the JACE by adjusting the leaving water set point. When load shedding was requested by GridMaster, the JACE would raise the set point, causing the chiller to cycle off. When load shedding was no longer necessary, the JACE would return the chiller water set point to its original value.

4.2.1 Field Test Architecture

The proposed control system applies to microgrids that can be islanded from the utility grid. An example of this is shown in Figure 30, where a section of a grid can be islanded by opening the microgrid isolation switch. This islanded section is shown with a yellow background.

There are three loads shown in the islanded section, although in a real-world installation there may be any number of loads. In this example, there is a low-priority load on the left that can be shed by opening the switch through which power passes to the building. Adjacent to that is the high-priority load that should not be shed (although it can be shed through a similar process). On the right is the load managed through OpenADR communication.

This microgrid also has three generators. In practice there may be more or fewer generators, along with energy storage and renewable sources. The power sources would ideally be sized to provide coverage of the critical microgrid loads, but it is a trade-off between energy security and the capital cost of installing assets that may only be used infrequently. Often, the local generation is insufficient to cover the entire load of the microgrid; consequently, load shedding will be necessary. In addition, power sources may not start when expected, may have an unanticipated fault, or may need to be taken offline for maintenance. When power sources become unavailable for any reason, a need for load shedding is usually triggered, sometimes by a significant amount.

The bottom row in Figure 30 shows the controller components. The S&C GridMaster controller has been successfully employed in DOD microgrids across the globe and was used in this demonstration. The Tridium BMS is used to control the heating, ventilating, and air conditioning (HVAC), lighting, and other loads in the building. The OpenADR VTN interacts with the Tridium controller to implement and verify strategic load shedding.

Forcing a DOD system into island mode for extended periods to do a proof of concept is both inconvenient and impractical. For this reason, the demonstration used the architecture shown in Figure 31. This architecture is a combination of simulation and actual loads, in which the actual loads are all low priority but some are configured as high priority in the controller. The simulation of the balance of the system runs on the GridMaster hardware as shown.

Normally, the GridMaster controller would communicate with the Tridium controller to obtain information about the system, as shown in Figure 30. However, for this demonstration it was only necessary to have GridMaster communicate with the OpenADR VTN. A temporary Niagara JACE controller replaced the existing Tridium controller algorithms and was programmed to respond to the load shed commands from the VTN.

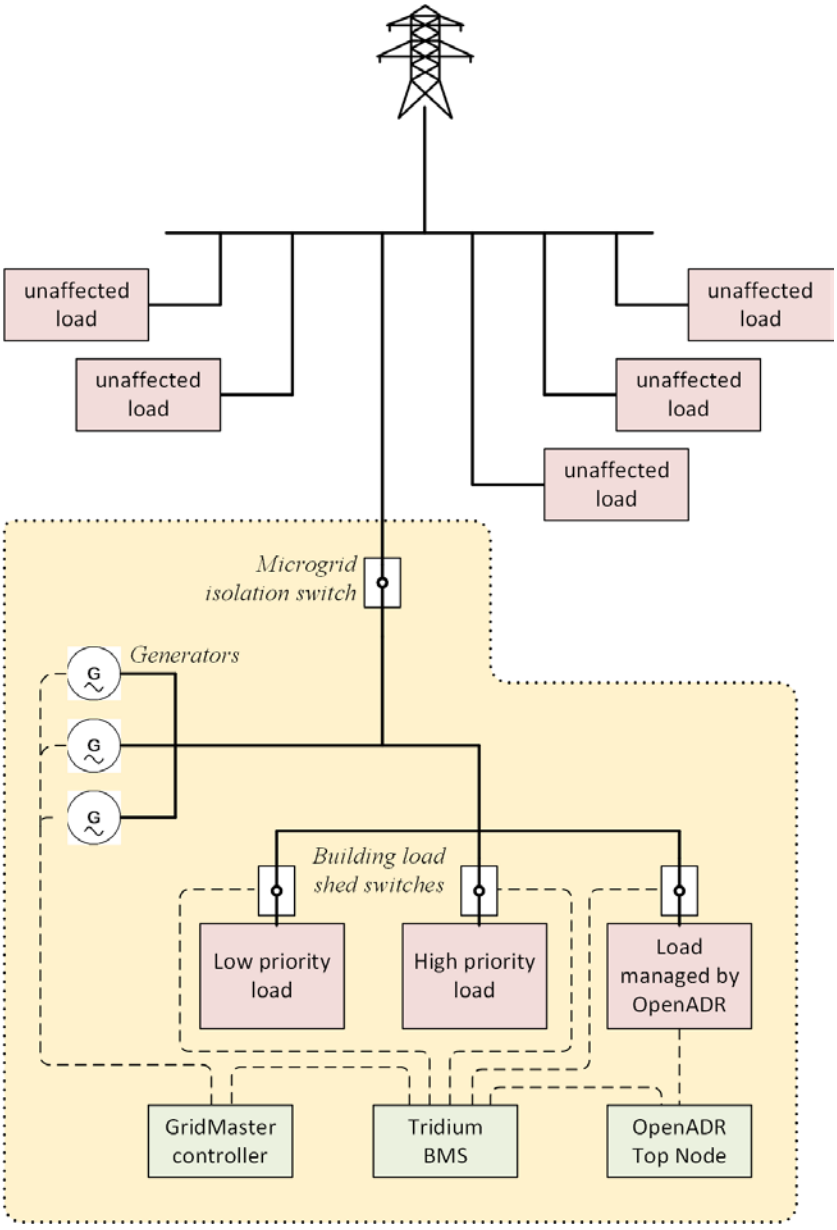

Figure 30. Idealized network topology for demonstration

The topology shown in Figure 31 indicates that the Tridium controller is communicating with several low-priority loads in the actual physical building. The actual nature of these loads does not matter from the perspective of GridMaster. The important aspect is to have the building loads respond to OpenADR commands within a reasonable time. Lighting panels and plug loads should respond almost immediately, whereas other loads (e.g., HVAC) may take several minutes to

respond. These response times are expected, although it is desired to have a full response in less than 5 min. That is, if a process—such as changing an HVAC set point—takes more than about 5 min to respond, then it is not a viable sheddable load for the purposes of this research.

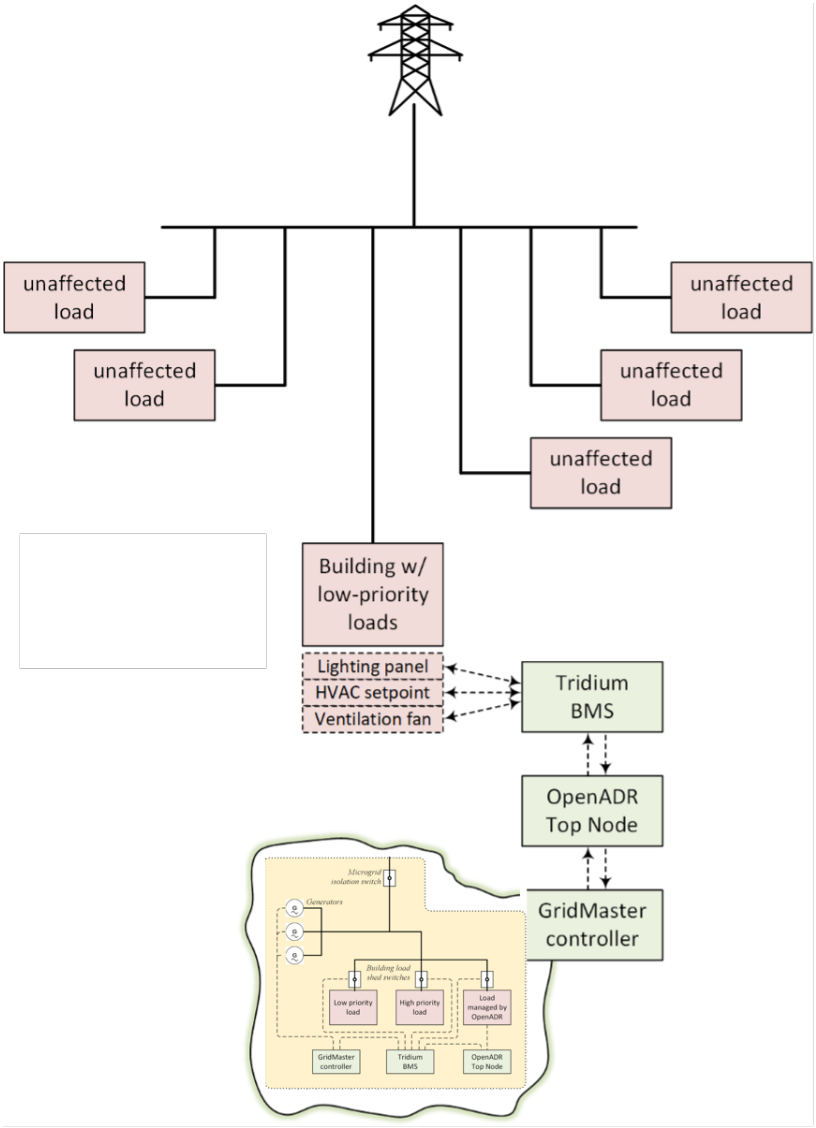

Figure 31. Proposed actual network topology for demonstration

4.2.2 Hardware Installation

To facilitate the tests, a remote-access system was installed along with the temporary JACE controller and the GridMaster controller. This allowed research personnel to remotely perform the tests and to run cybersecurity scans on the system. Figure 32 shows a one-line communication diagram of the various components used for the remote access. The SonicWall virtual private network (VPN) appliance ensured that the remote access was secure.

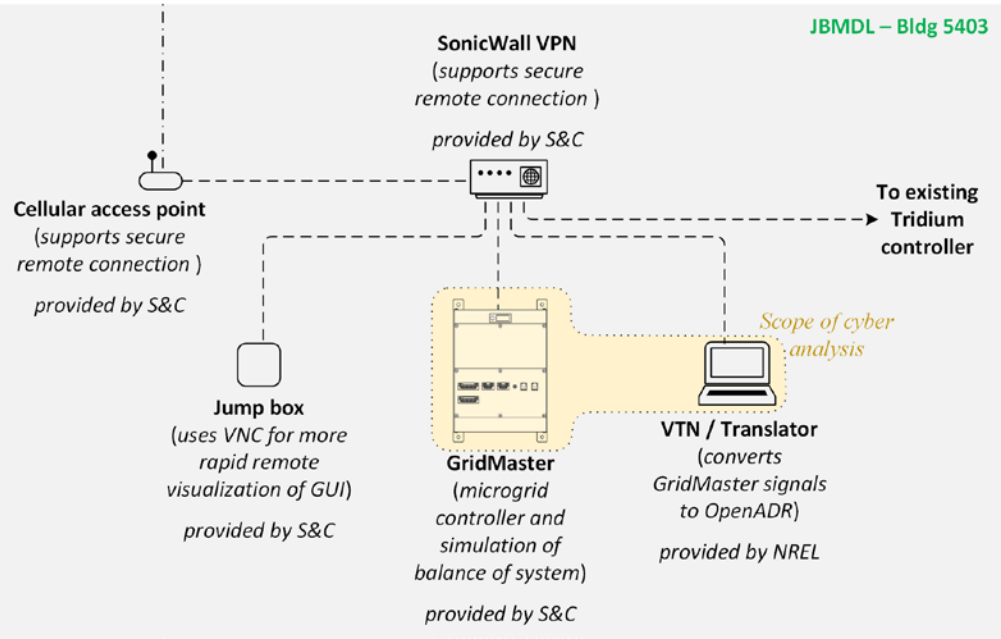

Figure 32. Hardware installed for control and remote system access

4.2.3 Hardware/Software Cost

For a facility with an existing MGC and a BMS supporting OpenADR communication protocol, the two main costs for utilizing this technology are the OpenADR server and software upgrades. Although OpenADR is open source, developing a VTN that issues load shedding signal could be time-consuming. For this project, we used a commercial OpenADR server from Grid Fabric, which cost $35,000.

The cost of the software upgrade depends on the complexity of the load flexibility strategy, ranging from simple upgrades that could be performed by facility managers at the site to complex updates that require control vendors. Determination of the software upgrades also require a thorough analysis of the energy-saving potential of assets controlled by the BMS.

In general, the software upgrade comprises two major updates. The first one is configuring the BMS to receive and distribute load shedding request from the OpenADR server to individual assets based on their load flexibility/shedding potential. The second upgrade requires updating the individual asset control system to those requests.

4.2.4 Field Test Results

The field tests were run at the facility under test (FUT), that is, Building 5403 at JB MDL. The simulated microgrid is shown in Figure 33, in which Load-01 represents the FUT, and the other equipment is simulated. The power consumption of the FUT was reported to the MGC by the JACE and was subsequently passed into the simulation. Most of the figures in the following section show the FUT power as reported by the simulation, which is taken from the power readings on the switch (SW-01) feeding power to that facility.

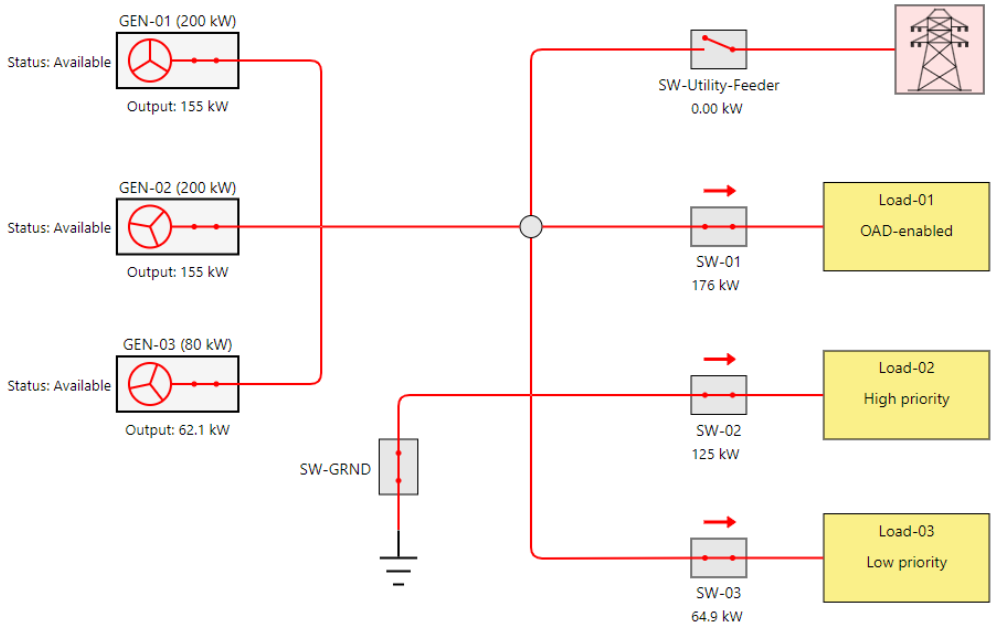

Figure 33. One-line diagram of simulated microgrid

Load-02 is considered a high-priority load and is not allowed to shed except under extreme conditions. Load-03 is designated as a low-priority load and can be shed through traditional means, that is, opening SW-03 to interrupt power to the entire building.

Approximately a month before the testing, it was discovered the only load in Building 5403 that could be controlled by the OpenADR system was the main chiller. This chiller has six stages with a total of six compressors, each dispatched separately according to the demand for cooling. Each stage of the chiller needs about 25 kW and was the dominant load in the building where the other loads totaled about 40 kW. There was concern the chiller load was large enough to cause undue oscillations in the load shedding process. It was also determined the whole-building power meter was nonfunctional.

These issues were addressed by creating a virtual whole-building power data stream equal to the sum of readings from several submeters and applying a multiplication factor of 3 to all the submeter data except for the chiller. The building loads are then more equivalent with those used during the validation tests at the ESIF lab at NREL, and the impact of the chiller load on the overall building load is diminished.

The JACE controller was originally programmed to shed the chiller load by sending a “disable” signal to the chiller controller. However, during initial commissioning, it was determined this was not a reliable method to manage the chiller load due to either large latency or unexpected chiller response. The load shed method was then changed to alter the chiller set point to use a value of 45°F leaving water temperature when the chiller was enabled and a value of 60°F to disable the chiller.

The chilled water served both a dedicated outdoor air system as well as recirculated air coils in the fan coil units throughout the building. During the time of these tests, the dedicated outdoor air

system fire alarms were in fault, rendering the outdoor air coils inoperative and consequently preventing the supply fan from running. All the loads on the chiller were therefore from the fan coil units, and because the building was unoccupied, the chiller was operating at a relatively low load. This introduced intermittency to the chiller operation, where it was cycling up to 10 times per hour. The tests were mostly run on humid days when the zone latent loads were high enough to allow the chiller to run continuously.

4.2.4.1 Field Test Case 1: Transition From Grid-Tied to Island Mode

The first test performed involved transitioning the microgrid from grid-tied mode to island mode when a power source was unavailable. The starting conditions had the simulated microgrid in grid-tied mode with all generation sources off and no shed loads. The smaller generator was then marked as unavailable—that is, it could not be dispatched to provide power to the microgrid. The user then initiated the process to put the microgrid into island mode through a closed transition. This process incorporates the built-in GridMaster islanding functionality with a custom algorithm that first determines if the system loads can be handled by the available generation sources. If not, it sets the allowed load for the FUT lower than the current load. Once the loads drop below the allowed threshold for islanding (80% of the available spinning capacity), the custom process then invokes the built-in closed transition process.

The results from this test are shown in Figure 34. In this figure, the purple line shows the baseline FUT load reported by the JACE controller, and the green line shows the measured load. The blue line shows the allowable power limit sent from the MGC to the JACE. This allowable limit is calculated based on the overall system load and the capacity and status of the available power sources on the grid. The red line in the graph is the power flow through the point of common coupling for the microgrid. The scale for the red line is on the right side of the graph.

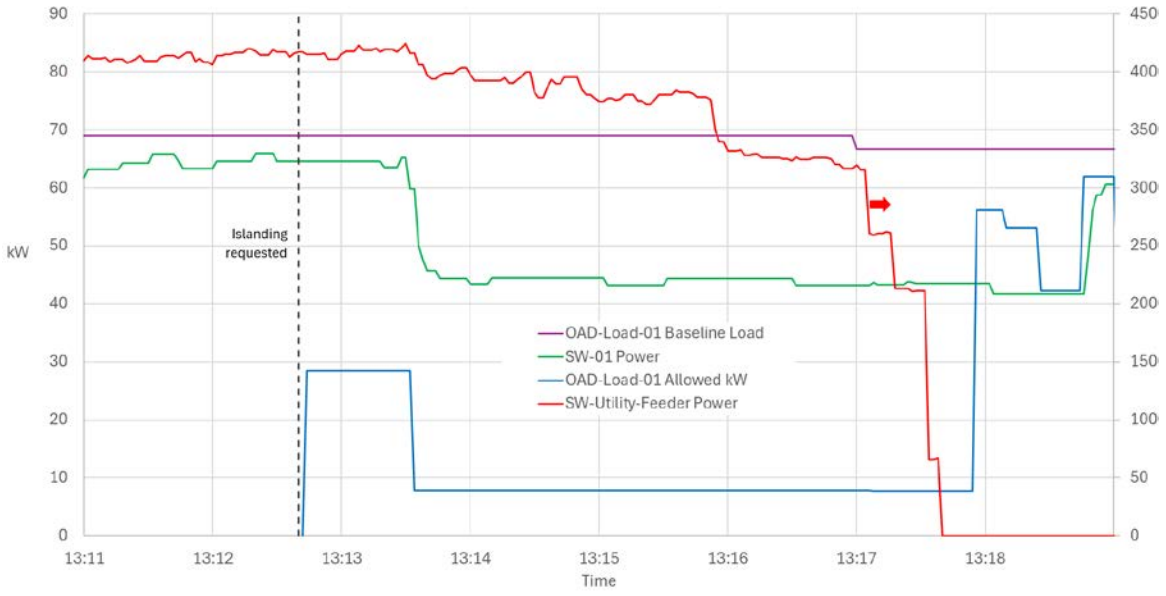

Figure 34. Field Test 1 results from actual equipment

The request to enter island mode was sent at 13:12:40, as shown by the dotted vertical line. The MGC immediately determined the overall microgrid load (>400 kW) exceeded the threshold for performing the transition and set the FUT allowable power level to about 30 kW. The chiller set point was modified accordingly, and approximately a minute later, the chiller shut off and the FUT demand decreased. The microgrid power demand remained greater than the threshold until about 13:18, at which point additional reduction in load was seen, and the islanding process was initiated. The generators were started and ramped up to drive the power flow across the power of common coupling close to zero, at which point the breaker was opened and the system was in island mode.

The original test plan assumed the generator that had been marked unavailable would be once again marked available so that the shed loads could be restored. In this test, the overall loads were low enough that it was not necessary because the full microgrid loads could be satisfied in island mode without the smaller generator. Once the system was islanded (after the utility feeder power went to zero), the allowed power calculated by the MGC was adjusted based on the observed load and the available spinning capacity. In this case, the overall loads were low enough the allowed power was greater than the existing building load (plus the hysteresis, not shown), and the chiller was reenabled and turned on around 13:18:40. This is, in effect, the same result that would have been achieved had the generator been marked as available.

The success metrics for this test were that (1) the microgrid islanding occurred within 5 min of the user requesting closed transition and (2) the loads were restored within 3 min of the generation source being made available. The load shed occurred within a minute after the island request was made, and the islanding occurred within 5 min. The shed chiller load was restored within a minute after islanding, so within the context of this test, it is considered successful.

4.2.4.2 Field Test Case 2: Load Variations While Islanded

This test validates the behavior of the system when loads increase or decrease while the microgrid is islanded. As critical loads increased, we expected the allowable FUT power level would decrease until loads were shed. As the critical loads then decreased, the shed loads would be restored. The starting condition for this test had the system already islanded with all three generators available and no load shed in effect.

Two attempts were made to demonstrate the behavior of the system during the scenario described. Figure 35 shows the results from the first attempt. The purple line shows the baseline FUT load reported by the JACE controller, and the green line shows the measured load.

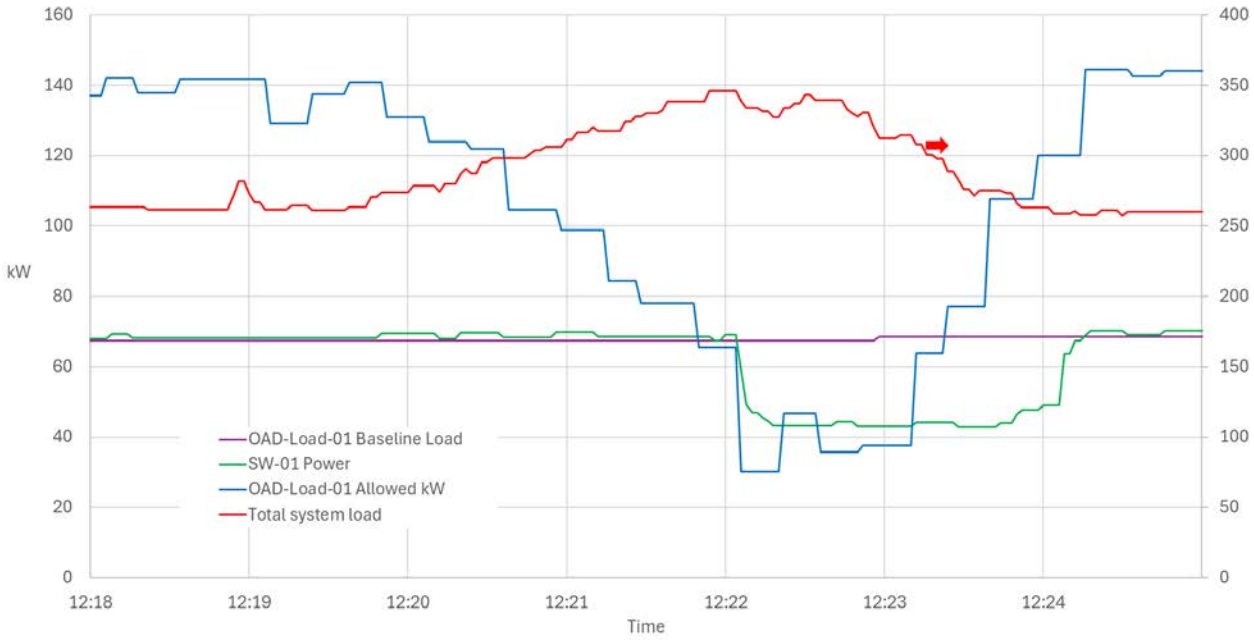

Figure 35. Field Test 2 results from actual equipment, first attempt

The blue line shows the allowable power limit sent from the MGC to the JACE. The red line in the graph represents the total power demand of the microgrid loads. The scale for the red line is on the right side of the graph. The loads on the simulated critical Load-02 slowly increased, starting around 12:19, until the allowed FUT load dropped below the measured load around 12:21:50. At that point, the MGC instructed the JACE to shed load, and the power to the building decreased approximately 20 s later.

The total system loads were then reduced to the point where the shed load could be restored. The allowable power exceeded the measured load around 12:23:25 and continued to increase until the allowed power plus the hysteresis was high enough for the JACE to restore the chiller shortly after 12:24.

The test was repeated, this time allowing the loads to increase until the simulated low-priority Load-03 loads were shed using traditional methods (i.e., opening a breaker to the entire building). The results are shown in Figure 36. The descriptions of the lines in this graph are the same as in Figure 35.

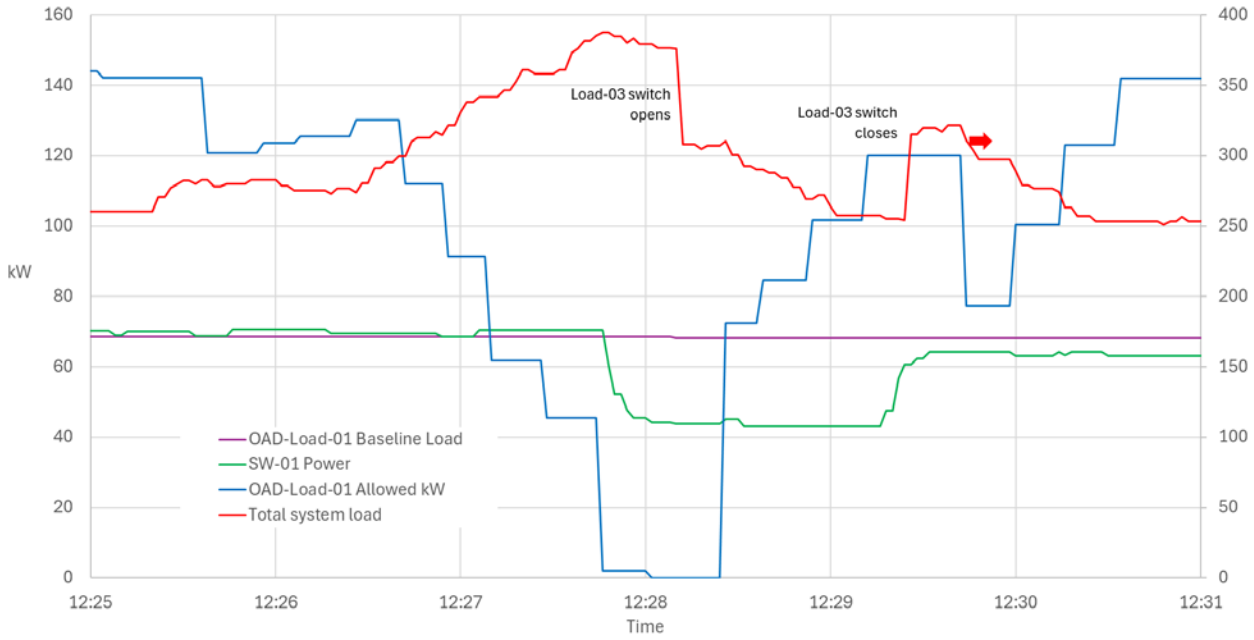

Figure 36. Field Test 2 results from actual equipment, second attempt

In this second attempt, the Load-02 demand was increased until the allowed FUT power decreased below the measured FUT load at 12:27:10. The chiller turned off, and the FUT load decreased approximately 45 s later. As the Load-02 demand continued to increase, the MGC eventually shed Load-03 by opening SW-03 at 12:28:10, as seen by the step decrease in the total system load. This resulted in the allowed power for the FUT increasing a few seconds later and the chiller loads being restored within about a minute. At the same time, the demand Load-02 slowly decreased until Load-03 was restored at 12:29:25, as shown by the step increase in the total system load.

The success metric for these tests was the ability of the controller to keep the total FUT load below the allowed threshold. Given that the response was as desired, this test was considered successful.

The test plan originally called for a third part to the test, in which high-priority loads in the FUT were increased to the point where the low-priority loads would be shed. The goal of this part of the test was to identify if sawtoothing of the loads could be avoided. Given that only one load of the FUT could be manipulated, it was not possible to execute this part of the test. However, given the development of routines specifically designed to avoid such oscillations described earlier, the researchers are confident this type of behavior can be avoided.

4.2.4.3 Field Test Case 3: Setting Generation Source Availability

This test validates the behavior of the system if generation sources are made unavailable while in islanded mode. As determined during the initial lab tests at the ESIF, OpenADR does not respond fast enough to handle sudden, unanticipated loss of a generation source. In this test, a generator is marked as unavailable in GridMaster, meaning it must be removed from service and cannot participate in the microgrid for long. The OpenADR system is then instructed to drop loads so that the now-unavailable generator can be turned off. The generator is then made available again. Once it turns back on, OpenADR should restore the loads that were shed.

The starting condition for this test is the system islanded with all generation sources available and no loads currently shed. The results from this test are shown in Figure 37, where the purple line represents the baseline FUT load reported by the JACE controller, the green line is the measured load, and the blue line shows the allowable power limit sent from the MGC to the JACE.

The smaller generator (GEN-03) was marked as unavailable at 13:00:05 in the MGC. This meant the controller considered this generator unable to participate in the microgrid and would attempt to shed load to the point at which the generator could be safely turned off. The allowed FUT power decreased immediately after the generator was marked unavailable and went below the measured load at 13:00:15. The chiller load was shed, and the FUT power decreased about a minute later. The generator was then set to available at 13:02:50 and eventually turned back on about 80 s later. This raised the allowed FUT power above the level where the hysteresis margin allowed the chiller to turn back on (at 13:05:25).

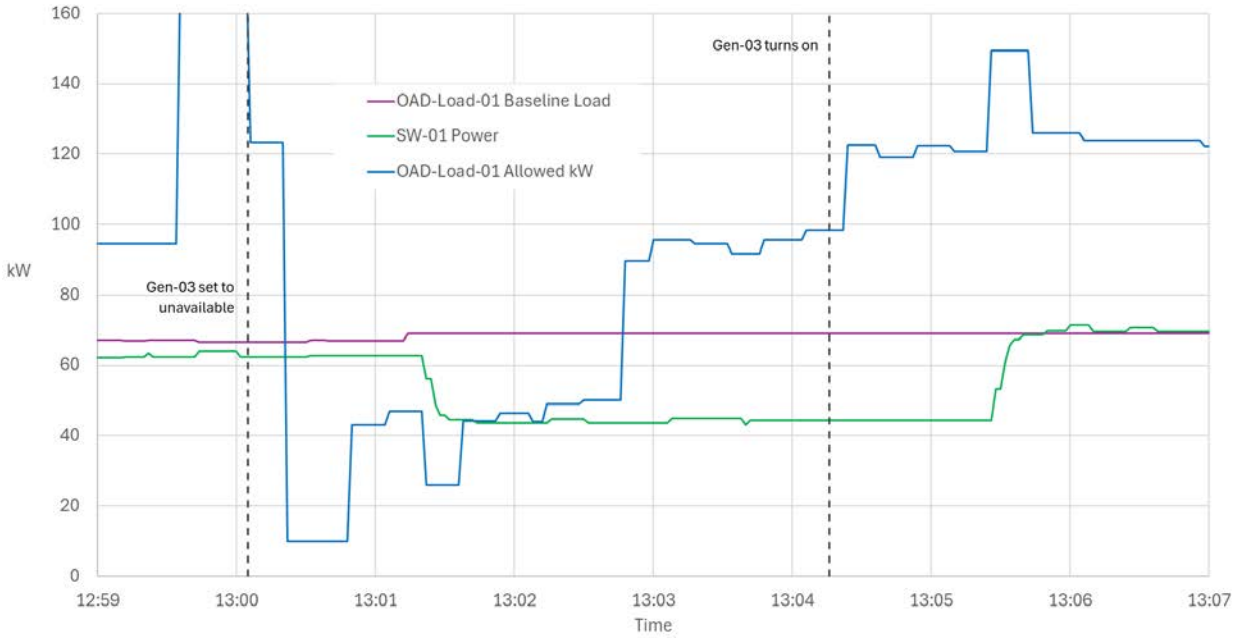

Figure 37. Field Test 3 results from actual equipment

The success metrics for this test considered the ability of the OpenADR system to reduce the FUT load to a level where the generation source (1) can be curtailed within 5 min of the generation source being marked as unavailable and (2) can restore load within 3 min of the generation source turning on. Both metrics were achieved, and this test was considered a success.

4.2.4.4 Field Test Case 4: Transition From Island to Grid-Tied Mode

This test validates the behavior of the control system once the microgrid has been reconnected to the utility. The system starts in island mode with increased critical loads and some noncritical loads shed. The user requests the system to reconnect to the grid, and GridMaster transitions the system to grid-tied and indicates to VTN that there is no need for load shedding. Once this happens, OpenADR should instruct the Tridium controller to restore all loads that have been shed.

The starting conditions for this test are the microgrid in island mode and the critical loads in Load-02 increased to the point where the FUT chiller load has been shed. The results from this test are shown in Figure 38, where the purple line represents the baseline FUT load reported by the Tridium controller, the green line shows the measured load, and the blue line is the allowable power limit sent from the MGC to the Tridium controller. The red line in the graph is the power flow through the point of common coupling and shows when the microgrid reconnects to the utility. The scale for the red line is on the right side of the graph.

The simulated Load-02 was increased until the allowed FUT load dropped below the measured load (at 13:08:10) and the chiller shut off and decreased the FUT load 20 s later. The MGC was then instructed to perform a transition back to grid-tied mode at 13:08:55, and the reconnection process completed about 40 s later (as shown by the utility feed power increasing in steps as the utility breaker was closed and the generators were sequentially turned off). Once the system was grid-tied, the allowed power was set to −1, indicating to the Tridium controller it is allowed to restore all loads that were shed. The chiller was restored at 13:10:50.

The metric for success of this test was to have all loads restored within 3 min of the transition to grid-tied mode. This metric was achieved, and the test was considered successful.

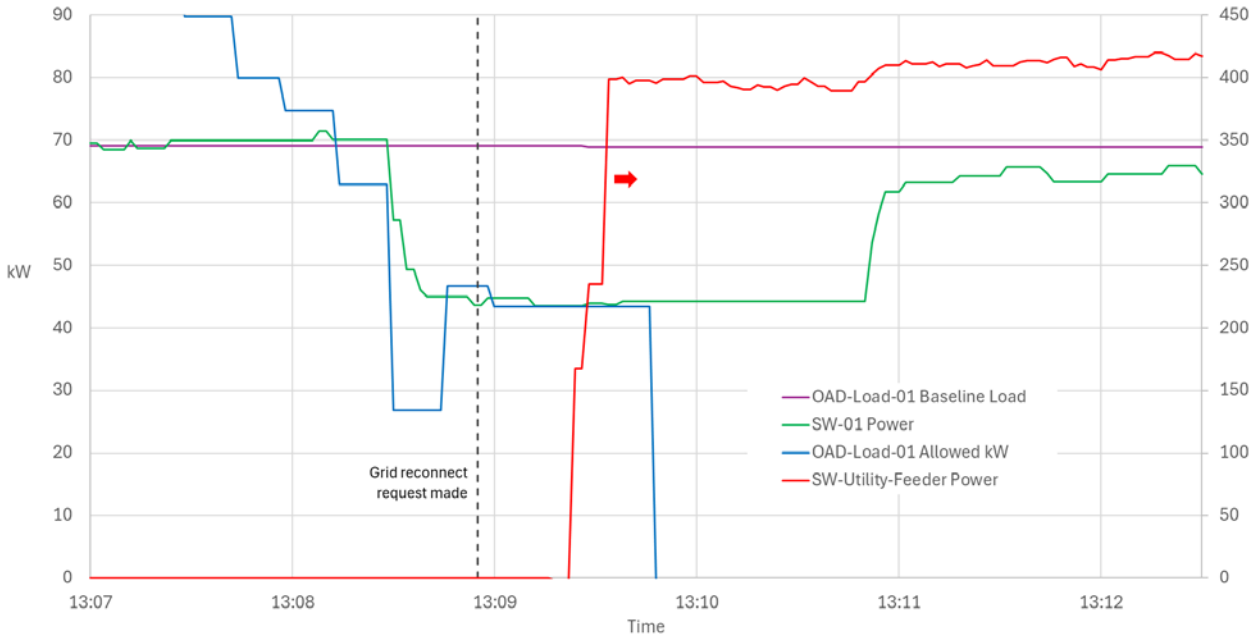

Figure 38. Field Test 4 results from actual equipment

4.2.4.5 Field Test Case 5: Round-Trip Message Timing Tests

During the course of these tests, the question arose as to how fast the OpenADR solution could be used to perform load shed. This is an issue that arises from the need to sometimes perform rapid load shedding in response to a fault in the system or the unexpected and sudden loss of a power source. Several tests were conducted to see how fast the load shed could reasonably be expected to occur. Figure 39 shows the typical results from these latency tests. The purple line shows the baseline FUT load reported by the Tridium controller, the green line represents the measured load,

and the blue line is the allowable power limit sent from the MGC to the Tridium controller. The red line shows the total system load scaled on the right side of the graph.

In the previous graphs, the Load-01 measured load was that reported from the simulation after the directly measured FUT power was passed into the simulation. In this graph, the green line shows the directly measured building load and avoids any latency introduced by the reporting of the simulation values.

In this test, the Load-02 demand was increased until the allowable FUT load dropped below the measured load. As shown in the graph, once this occurred, it took 25 s for the various calculations to be made and commands implemented to curtail the chiller. This falls in line with the previous tests described in this report, where response times for the chiller load shed are in the range of 15 to 45 s. This intuitively makes sense because there are many different components and calculations in the calculation chain. For example, in this case, the data path is:

Total load measurement → Allowable load calculation → Updating the Tridium controller with new allowable load and hysteresis value → Comparison of allowable load with measured building load → Command of set point change to chiller → Chiller internal control response → Total load measurement and reporting of new values.

Each of these steps will take a few seconds and may introduce additional latency, depending on where each component is in the scanning and calculation cycle.

Figure 39. Typical results from response latency test

4.3 CYBERSECURITY EVALUATION

In September 2023, a cybersecurity analysis was performed on hardware and network communications used when the project's specific OpenADR implementation was tested in a simulated environment at the ESIF facility. A second scan was performed in June 2024 during the demonstration of the system at JB MDL. Both the initial testing and on-site demonstration used

the same protocols, a similar communications scheme, and similar hardware platforms. The primary focus of the demonstration was to assess the functionality of the system, so extensive cybersecurity hardening was not performed, as it would likely complicate the initial efforts to get the system communicating and working.

As shown in Figure 32 and highlighted in Figure 40, the scope of the analysis was limited to the MGC and the associated protocol translator.

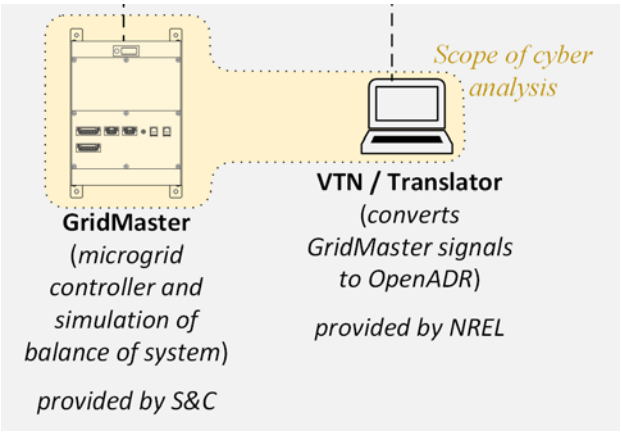

Figure 40. Hardware installed for control and remote system access

The following sections describe the results of the scans from the system as implemented at JB MDL; discuss the available security profile of the hardware, software, and protocols used; and make recommendations concerning likely requirements and cybersecurity enhancements applicable to a more permanent installation. The following methods and scanning tools were used to evaluate the cybersecurity stance of the implemented solution:

- Network Mapper (Nmap) and Nessus scans were executed against the Modbus translator, which is a special custom-designed appliance to translate ADR communications to Modbus/TCP accessible to the GridMaster. The same scans were performed against the GridMaster controller.
- Wireshark traces were performed on the local area network (LAN) to analyze dataflows among the tested devices.

4.3.1 OpenADR Protocol as Implemented

The use of OpenADR in conjunction with local BMS offers functional benefits for refinement and granularity in managing building loads and offers the security benefits of Transport Layer Security (TLS) when OpenADR's secure version is implemented. Due to complexities in certificate management and initial configuration, the HTTP (rather than the secure HTTPS) version of the protocol was used. The HTTP version offers very few security features and has significant vulnerabilities. For example, although a password is used for authentication to the VTN, the password is transmitted in plain text and can be obtained by monitoring the network. As shown in Figure 41, the password was obtained by inserting a network tap in one of the Ethernet cables in the data path between VEN and the translator and monitoring the communications using Wireshark.

Similar methods could be used to obtain information about the electrical loads on the site. Additionally, the use of HTTP could allow another device on the LAN to impersonate either the VTN or the VEN, as the protocol does not provide for meaningful device authentication. The lack of authentication and encryption could allow a malicious actor to gather information about activities on the site as well as interfere with operations, potentially causing disruptions to critical loads and potentially allowing damage to electrical infrastructure on site.

For these reasons, it is recommended to implement the TLS version of OpenADR, with the attendant public key infrastructure trust chain. Doing so would allow for strong encryption and authentication for ADR communications and devices and would protect the system from intrusion, even if the intervening network was compromised. For the anticipated use case, where OpenADR is used to communicate with multiple BMS, using a secure protocol across a campus would be preferable.

Use of HTTP will require securing the network connecting these devices, physically and/or with appropriate technical network controls, which will be more challenging as the OpenADR network expands.

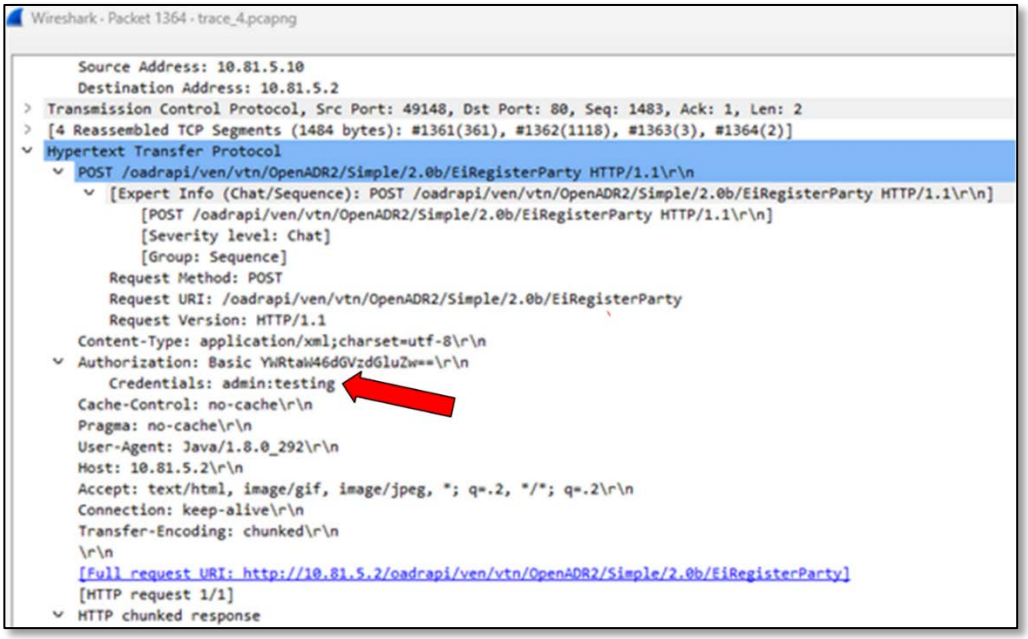

Figure 41. Wireshark screen capture showing plain text transmission of credentials

The network created to support these devices would need to be designed to protect the sensitive communications and carefully restrict access. The likely approach would be to create a walled garden, with the innermost section of the network containing the sensitive operational technology devices, with restricted access allowed only to external information technology devices (e.g., HMI) that require it. The security controls to implement this approach would likely include the following:

- Physical security to protect access to the switches, routers, cabling, and devices themselves.
- Logical controls such as access control lists applied at switchports and routers, or virtual LANs.

- Careful restrictions at the boundaries of the network, which may connect to other networks on the system. In the case of connections located at remote buildings, care will be required to ensure these field points of connection cannot be used to gain unnecessary access to the inner network or other remote locations.
- Careful consideration of how nonoperational technology devices (such as HMI computers and maintenance computers) connect to the inner network.

The Purdue five-level diagram is used at DOD sites to visualize a hierarchy of device types contextualized in a wider system. Devices are divided among five levels, described in Table 4. Devices are stratified in the Purdue diagram based on the communications media the device uses, and by the extent of the device's communication to other networks.

Table 4. Purdue Security Model for Industrial Control Systems

Purdue Level	Description
1.1.1.1 5	1.1.1.2 Interfaces to "external" networks
1.1.1.3 4	1.1.1.4 Sitewide Internet Protocol (IP) network used for the control system, along with front-end servers and workstations (desktops and laptops)
1.1.1.5 3	1.1.1.6 Connection between the field control system IP network at Level 2 and the Level 4 IP network
1.1.1.7 2	1.1.1.8 Networked controllers on an IP network
1.1.1.9 1	1.1.1.10 Networked controllers not on an IP network
1.1.1.11 0	1.1.1.12 Non-networked devices which communicate using analog signals

If deployed at a DOD facility, the Purdue diagram for the system at NREL would likely be represented as shown in Figure 42. Such a diagram would likely have to be created for the various cybersecurity documentation necessary for the installation to receive an authority to operate (ATO). In this diagram, the HMI computers and log aggregator (syslog) reside at Level 4 along with other devices used for human interaction. These use predominantly secure protocols to connect with the Level 2 devices (secure protocols are shown in green, insecure in red). Their access to the Level 2 (sensitive devices) would be restricted by the switches or router implemented in Level 3.

The Level 2 components that use insecure protocols for device communications would reside in a protected virtual LAN. The BACNET connection to the buildings (orange area) would be the most problematic to secure. If the HTTPS OpenADR protocol were implemented, it might be possible to locate the Tridium(s) at the remote sites, as HTTPS would be more easily secured through the intervening distance.

Syslog communications can be secured via TLS. However, Tridium/Niagara does not seem to offer that capability. The network infrastructure (e.g., managed switches) can also be configured to send logs to the syslog server. This is useful for monitoring and reviewing network connections/disconnections and activity through the generated logs.

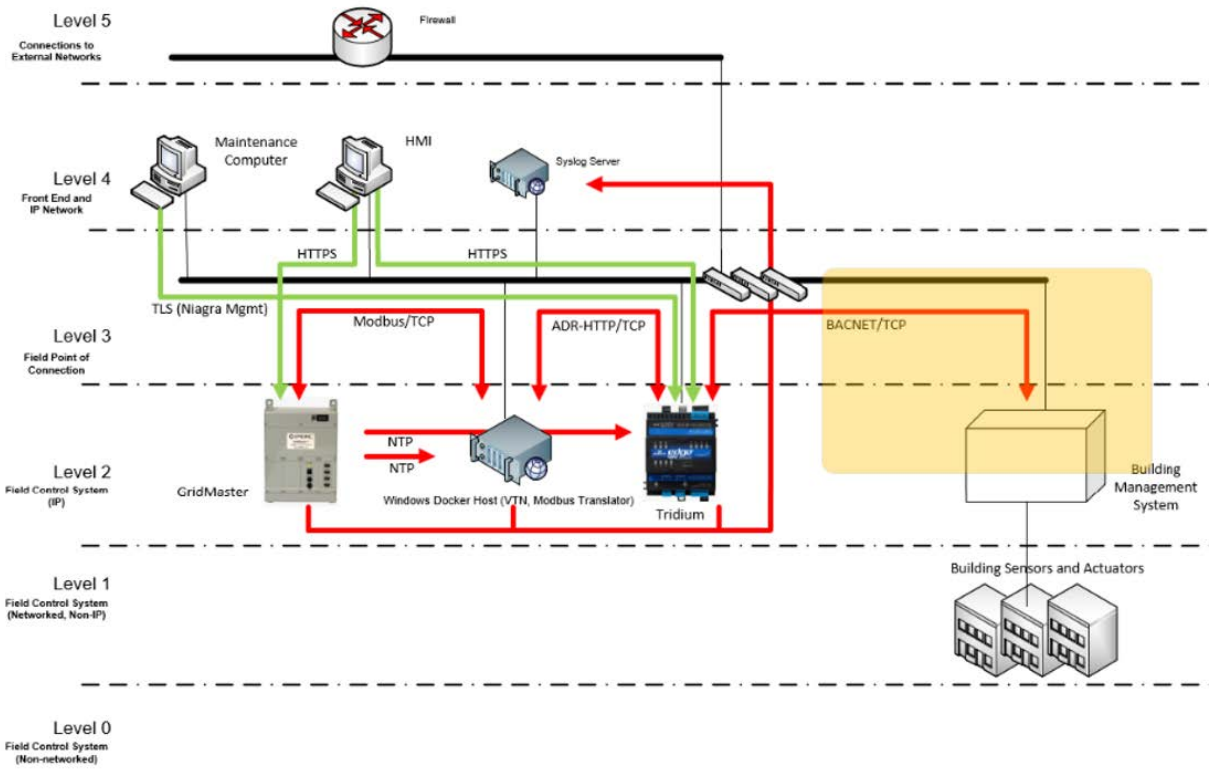

Figure 42. Representative Purdue diagram with OpenADR components

There were some transmission errors and warnings detected during the cybersecurity test. These included "spurious retransmissions" and "acknowledgements for unseen segments." Because the system was in a testing environment, some components of the system were simulated and shared the same host and IP addresses, which likely contributed to the errors. There was possibly a failed Ethernet cable involved. In an actual deployment, it will be necessary for communication paths to be consistent and predictable, with minimal retransmissions and other irregularities. Restricting IP addresses and network interface cards (NICs) to a particular function would aid in securing the network. For example, the IP address of the Modbus Translator would not be used to initiate an HTTPS session with the GridMaster UI webpage. It would be desirable to assign the Modbus translator to a dedicated host, which would not be used for human interaction with the system. This would allow the host to be secured via security controls specific to its function and would assist in monitoring the device for appropriate behavior.

The cybersecurity assessment did not test or evaluate the ability of the system to recover from errors in a safe and secure state. In a real-world implementation, it will be necessary to consider various modes of failure, how recovery will take place, and what effect various failure modes can have on the mission of the DOD location. Redundant network paths that use Spanning Tree to prevent loops may be necessary to maintain operation in the event of a trunk link or switch failure.

4.3.2 Modbus Translator

The Modbus Translator is an application that serves as an interface between the GridMaster controller and the OpenADR VTN. This piece of equipment (or something similar) will be

necessary in future deployments to convert messages from equipment that cannot send arbitrary HTTP/S requests in the format the VTN expects. The Translator runs in a docker container on a Windows host and offers both an HTTP server (for OpenADR use by the VEN) and a Modbus server (for use by the GridMaster controller). In a permanent installation, the cybersecurity posture of both the Windows host and the container must be considered. As evidenced by the Nmap scan of the Modbus translator, a substantial bit of information can be obtained over the network concerning the libraries and structure of the docker container itself and the Windows host it is operating on. Examples of information obtained through Nmap scanning of the on-site demonstration at JB MDL are shown in Figure 43.

From the network scans performed at JB MDL, it is possible to identify that an NREL-associated computer is running a VTN server. The use of ruby libraries can be identified as underlying components of the VTN appliance. The LANDesk and EnCase services are used for information technology management and would likely not be used in a permanent installation. Limiting the information provided freely by network hosts helps reduce the information available to a potential attacker.

Reducing listening ports by disabling them and restricting the IP addresses that can access the necessary ports (accomplished via a host-based firewall or by managed network hardware) are effective hardening actions that reduce the vulnerability of the network to cyberattacks.

```
Target_Name: NREL_NT
| NetBIOS_Domain_Name: NREL_NT
| NetBIOS_Computer_Name: SDESK-34402S
| DNS_Domain_Name: nrel.gov
| DNS_Computer_Name: SDESK-34402s.nrel.gov
| DNS_Tree_Name: nrel.gov
|_ Product_Version: 10.0.2262
```

```
| /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/actionpack-6.0.5.1/lib/action_dispatch/middleware/debug_exceptions.rb:36:in `call'
| /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/actionpack-6.0.5.1/lib/action_dispatch/middleware/show_exceptions.rb:33:in `call'
| /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/railties-6.0.5.1/lib/rails/rack/logger.rb:37:in `call_app'
| /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/railties-6.0.5.1/lib/rails/rack/logger.rb:26:in `block in call'
| /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/activesupport-6.0.5.1/lib/active_support/tagged_logging.rb:80:in `block in tagged'
|  /canvas-oadr-vtn/bin/jruby-9.2.13.0/lib/ruby/gems/shared/gems/activesup
```

```
Discovered open port 135/tcp on 10.81.5.8
Discovered open port 3389/tcp on 10.81.5.8
Discovered open port 33354/tcp on 10.81.5.8
Discovered open port 9595/tcp on 10.81.5.8 LANDesk/Intel Management Agent
Discovered open port 4445/tcp on 10.81.5.8 EnCase Servlet
Discovered open port 3000/tcp on 10.81.5.8 /canvas-oadr-vtn/bin/jruby-9.2
Discovered open port 9594/tcp on 10.81.5.8 LANDesk/Intel Management Agent
Discovered open port 9593/tcp on 10.81.5.8
```

Figure 43. Nmap scanning examples

If the projected use case includes providing OpenADR communications to multiple VENs across the campus, the security requirements for the server would likely be more stringent, as the device would serve multiple subsystems spanning a larger geographical area.

4.3.3 GridMaster Microgrid Controller

The GridMaster is a controller used to manage loads, sources, and situations on a microgrid. It provides human interaction by offering a TLS webserver. If other GridMasters are used in the system, these communications are also secured by TLS and use certificates, which can be generated specifically for the site. TLS is a secure protocol.

To communicate with microgrid end devices, the GridMaster uses supervisory control and data acquisition (SCADA) protocols; Modbus is used to communicate with the Modbus Translator. SCADA protocols such as Modbus are inherently insecure, and the network paths that support these protocols must be appropriately managed and protected.

The GridMaster used in the JB MDL demonstration was found to be listening on the ports, as shown in Figure 44.

```
Discovered open port 22/tcp SSH for maintenance
Discovered open port 3/tcp Used for GridSim simulation
Discovered open port 9001/tcp Used for GridSim simulation and database
Discovered open port 8443/tcp TLS connection from HMI
Discovered open port 32/tcp SSH for maintenance
Discovered open port 4/tcp Used for GridSim simulation
Discovered open port 123/udp NTP
```

Figure 44. Open GridMaster ports

The secure shell (SSH) ports (22 and 32) are considered secure and used only for GridMaster maintenance sessions by installation and maintenance engineers. Port 8443 is used to provide a TLS webserver connection to the HMI computer. Ports 3, 4, and 9001 are related to GridSim and would not be used in an actual deployment. The network time protocol is open, as the GridMaster can provide a GPS time source to other devices on the network.

In a permanent deployment, the GridMaster's internal firewall would be enabled and would restrict access to only the IP addresses allowed for the given services (e.g., only the HMI computer would be allowed access to the HMI webserver). The GridSim ports would not be exposed to the network. All open ports would be rate-limited to prevent attacks on DOD.

To secure management connections to the GridMaster over TLS, and to secure inter-GridMaster communications, it will be necessary to generate and install an appropriate TLS trust chain, a service that S&C provides for DOD GridMaster installations. This trust chain could also be incorporated into the OpenADR appliance, as well as the Tridium controller, which would facilitate device authentication throughout the system.

4.3.4 Virtual End Node (Tridium Station)

The Tridium station uses an OpenADR connection to the VTN node to read information related to the maximum allowable building load and write information concerning the available load shed and baseline load. The Tridium, in turn, takes steps to monitor status and implement desired control

changes to the buildings. The Tridium controller was not considered in the cyber evaluation, as it would likely be part of the building management enclave. The Tridium was tasked with monitoring and controlling numerous end devices in the building. The Tridium communicates with these end devices using BACnet; it uses a separate, dedicated physical NIC to communicate with the VTN and GridMaster. This serves to isolate a building's subsystem from the OpenADR communications.

If an authorization boundary is placed between the building's BMS (Tridium) and the OpenADR/microgrid equipment, special attention will be needed to restrict traffic across this boundary. In practice this would likely mean installing a firewall, restrictive access control lists, or other types of network control at this boundary, and allowing only the required communications (OpenADR) to traverse. The Tridium management should not be exposed to Modbus Translator, and the management services on the translator should not be accessible from the building management networks.

4.3.5 General Cybersecurity Recommendations

Based on the findings from the lab and field test, the following recommendations are made with respect to the proposed solution:

- Although not tested in this iteration, using the TLS version of the ADR protocol would greatly benefit cybersecurity and reduce the stringency of required network controls. Applying encryption would not cause a substantial increase in bandwidth usage or latency in the system and would allow more flexibility in the design of the network, as less protection and isolation would be required for the encrypted version. The OpenADR protocol is not a high-speed protocol with demanding latency requirements, so use of TLS would likely not have a significant impact.
- All devices on the system will need to have unnecessary ports and services removed or blocked, to reduce the attack surface present on the network. This includes simulation-related listening ports on the GridMaster and any unnecessary Windows-related ports on the translator. This is especially applicable if an authorization boundary is placed between the translator and the BMS.
- The network should be designed to isolate and protect insecure SCADA protocols, such as Modbus, BACnet, and OpenADR, until a secure version of the OpenADR can be incorporated and used.
- For the GridMaster controller, the usual DOD hardening process with the latest Security Technical Implementation Guides requirements will have to be applied. A plan should be generated that addresses maintenance of the underlying operating system with the latest security patches.
- For the Modbus translator, the Windows platform should undergo the established Security Technical Implementation Guides process, and a method should be established for maintaining the platform and docker with appropriate updates, as both present potential vulnerabilities.
- Procedures should be established for backing up and restoring all the necessary equipment and the ability to reset them to known good configurations, in the event of malfunction or a cyber incident, in which tampering can be suspected.

- A centralized log server should be procured, or manual methods should be established for retrieving log information from the GridMaster, Tridium, and Modbus Translator. These devices should be configured to log events of cybersecurity interest, such as successful/unsuccessful logins, connection attempts, user changes, and software/configuration changes.
- Consideration is needed to establish appropriate user accounts and access, following principles of least privilege, so that devices can be monitored and maintained as necessary by on-site personnel.
- If an authorization boundary is placed between the ADR equipment (translator) and the BMS (Tridium) systems, special care is needed to limit the protocols and end points that can traverse the boundary.

5.0 IMPLEMENTATION CHALLENGES

Although this technology has significant potential to provide selective control of building assets during power shortages and prevent total power loss in noncritical facilities, there are implementation challenges that could affect its wider application. The main challenges are discussed in this section.

5.1 OpenADR PROTOCOL ADOPTION AND COMPATIBILITY

The OpenADR protocol is relatively new, and its widespread, native integration across the spectrum of BMS has not yet been achieved. Many legacy and even mid-range BMS products currently lack this native support. Consequently, benefiting from this technology requires facilities to either possess a compatible, modern BMS or invest in middleware, gateways, or software customizations to bridge the communication gap with their existing infrastructure. This limited native compatibility presents a significant barrier to broad adoption, adding complexity and cost to the integration process and narrowing the pool of viable facilities. It is also worth noting the original intent of OpenADR is to provide a standard communication protocol between a utility and their customers to facilitate consistent methods for implementing demand management. The use of OpenADR for this project represents a novel application of the standard with its own challenges as described in the report.

5.2 QUANTIFYING LOAD FLEXIBILITY POTENTIAL

The actual benefits derived from this technology are directly proportional to a facility's ability to accurately identify and quantify the load flexibility potential of its various assets. This assessment, which includes understanding the operational interdependency of critical process requirements, occupant comfort thresholds, and the dynamic behavior of diverse equipment (e.g., HVAC systems, lighting, refrigeration, industrial machinery), can be challenging for complex systems found in microgrids. Without granular data and a precise understanding of what can be curtailed, for how long, and with what operational impact, facilities risk either underutilizing their flexibility or inadvertently disrupting critical operations, thereby undermining the perceived value of the solution.

5.3 COSTS ASSOCIATED WITH CONTROL SYSTEM UPGRADES

Implementing an OpenADR solution typically necessitates a multi-tiered approach to control system upgrades. This includes updates at the central BMS level to receive and interpret OpenADR signals, as well as modifications or additions to the control systems of individual assets (e.g., chillers, air handling units, lighting circuits, smart thermostats) to enable them to respond directly to load shedding requests. Facilities may find these substantial upfront costs difficult to justify, particularly if the payback period or quantifiable benefits are not immediately clear or sufficiently compelling to outweigh the initial investment.

In addition, implementation of the load management algorithms in the OpenADR top node requires understanding and collaboration with the facilities management team so that the various equipment can be operated in a useful and safe manner. The participation of facilities management can increase the overall cost.

5.4 LATENCY LIMITATIONS FOR RAPID RESPONSE SCENARIOS

As demonstrated in section 4.0, the inherent communication latency of the OpenADR protocol makes it unsuitable for scenarios demanding immediate few-second level responses. OpenADR is designed primarily for ADR signals that typically operate on timescales of minutes or more, aligning with utility-initiated events rather than instantaneous operational adjustments. Consequently, it cannot effectively address critical, fast-acting events such as a sudden loss of grid power, an unexpected surge in internal load, or the pre-shedding requirements for a black start operation. In these urgent situations, facilities must continue to rely on conventional, localized control mechanisms such as switchgear or protective relays. This limitation means this technology does not provide a holistic solution for all load management needs, and its inability to handle rapid, safety-critical responses might discourage facilities from investing, perceiving it as an incomplete or insufficient solution for their comprehensive operational demands.

5.5 NETWORK HARDENING AND CYBERSECURITY IMPLICATIONS

The integration of an OpenADR server inherently connects a BMS to external networks, significantly expanding its digital attack surface. This necessitates robust network hardening measures to protect against potential cyber incidents that could compromise energy efficiency, operational integrity, or even safety. Depending on the existing cybersecurity structure and the availability of skilled personnel and financial resources at a given facility, this critical security requirement can add a considerable burden in terms of initial investment and ongoing maintenance.

6.0 CONCLUSIONS AND IMPLICATIONS FOR FUTURE RESEARCH AND IMPLEMENTATION

The use of the OpenADR communication protocol to perform strategic load shedding in islanded microgrids was successfully demonstrated in simulated environments as well as in a control hardware-in-the-loop setting. For these tests, a commercial MGC was used with normal programming within the expected constraints of system capabilities. This means the solution did not require any specialized modification to the code base of the controller.

Tests were conducted that considered the load shedding required to achieve a closed transition into island mode and to account for changing loads and power source availability while islanded. Although the field tests only considered a single controllable load in the physical building, the algorithms have previously been shown to work with multiple smaller loads in a laboratory setup. It is expected the demonstrated solution could be a practical method for implementing strategic load shedding in similar projects.

The technology is based on the communication of five parameters: baseline load, load power, available load shed, allowable load, and hysteresis margin. The first three are transmitted from the VEN to the MGC, whereas the last two are transmitted from the MGC to the VEN. The load power indicates the current power use of the building and thus requires submetering the building under consideration. Baseline load is the estimated power use of the building if it was not in a shedding event. If there is no shedding request, the baseline load is the same as power load. Available load shed is the amount of load that could be shed by modulated controllable assets in the building. Determination of both baseline load and available load shed requires a good understanding of the system and involves modeling the assets' energy use. The complexity of the models depends on the type of assets under consideration.

Considering the latency of the round-trip communication path between the MGC and the different devices/assets involved in the load shedding processes, there are certain scenarios for which the demonstrated solutions are suitable and others for which they are not. When the microgrid is islanded, the method described in this report could be used to set target power demands for buildings within the microgrid. This would be done at the time of islanding and be based on the mission criticality of each building. A custom process would be created that compares the sum of capacities of all available power sources to the sum of all anticipated building loads. If there is insufficient generation, then OpenADR would be used to limit/shed building loads to manageable values. If a system is operating in islanded mode, the method could be used to accommodate normal variations in load and the need to remove a power source from operation for maintenance.

There are several scenarios for which OpenADR is not an appropriate mechanism to implement demand reduction. These are mostly due to the inherent latency of OpenADR response to a request. The scenarios include:

- **When necessary to decrease loads rapidly due to an unexpected loss of a power source:** The latency test during the field test indicated the response times for the building system to respond to the shedding request are in the range of 15 to 45 s. A generator can run for a few seconds over its rated capacity, depending on how far over its capacity the generator is operating. If the grid is already operating with an N+1 or N+2 strategy, then it

can survive the unexpected loss of a power source; if it is operating with an N strategy, then the delay of the OpenADR response is likely too slow to preserve the grid.

- **Attempting a black start that requires pre-shedding of loads prior to energizing the bus:** This scenario requires prior knowledge of what the building loads will be once the building is reenergized. During the black start, it is difficult to know the loads in the building unless the loads are schedule-based loads that are easy to determine.
- **Reducing demand on a microgrid due to a sudden and unexpected increase in load:** Although the grid may survive a sudden load increase if the step change is small enough, a large step change could require a rapid load shed (within a few seconds) to avoid overwhelming the spinning generation. This rapid load shed is not possible due to the inherent latency problem.

To ensure the proper functioning of the system and avoid any vulnerability of the system, cybersecurity hardenings are needed. Section 4.3.5 summarizes cybersecurity recommendations that need to be in place for implementing the technology.

7.0 REFERENCES

Benne, Kyle, Brian Ball, Willy Bernal Heredia, Dylan Cutler, Stephen Frank, Larry Brackney, Anya Petersen, Tobias Shapinsky, Nicholas Long, Tim Coleman, Alex Swindler, Eric Bonnema, and Lauren Klun. 2018. "Alfalfa [SWR 18-17]." Computer software. November 14, 2018. https://github.com/NREL/alfalfa. https://doi.org/10.11578/dc.20181205.2.

Deru, Michael, Kristin Field, Daniel Studer, Kyle Benne, Brent Griffith, Paul Torcellini, Bing Liu, et al. 2011. *U.S. Department of Energy Commercial Reference Building Models of the National Building Stock*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-46861. https://www.nrel.gov/docs/fy11osti/46861.pdf.

OpenADR Alliance. 2012. "Berkeley Labs Advances Open ADR." Accessed July 31, 2024. https://openadr.memberclicks.net/index.php?option=com_content&view=article&id=70:berkeley-labs-advances-open-adr&catid=29:news.

OpenADR Alliance. 2023. "OpenADR Momentum & Ford Joins Board of Directors." Accessed July 31, 2024. https://www.openadr.org/index.php?option=com_content&view=article&id=208:openadr-momentum---ford-joins-board-of-directors&catid=21:press-releases&Itemid=121.

APPENDIX A SCIENTIFIC/TECHNICAL PUBLICATION

Curtiss, Peter, Korbaga Woldekidan, Donald Price, Tobias Shapinsky, Scott Roberts, and John Gehring. 2024. *Integrating a Microgrid Controller With a Local OpenADR Server*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-90249. https://docs.nrel.gov/docs/fy25osti/90249.pdf.