Integrating a Microgrid Controller with a Local OpenADR Server

Peter Curtiss,[1] Korbaga Woldekidan,[2] Donald Price,[2] Tobias Shapinsky,[2] Scott Roberts,[1] and John Gehring[1]

1 S&C Electric Company
2 National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5500-90249
October 2024

Integrating a Microgrid Controller with a Local OpenADR Server

Peter Curtiss,[1] Korbaga Woldekidan,[2] Donald Price,[2]
Tobias Shapinsky,[2] Scott Roberts,[1] and John Gehring[1]

1 S&C Electric Company
2 National Renewable Energy Laboratory

Suggested Citation
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://www.nrel.gov/docs/fy25osti/90249.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/5500-90249
October 2024

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) 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.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

List of Acronyms

ADR	automated demand response
BMS	building management system
HVAC	heating, ventilation, and air conditioning
MGC	microgrid controller
OAD	OpenADR device
VEN	virtual end node
VTN	virtual top node

Executive Summary

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.

The report also includes test plans for evaluating the performance of the software upgrade under various operating conditions typical of microgrid applications.

Table of Contents

1 Project Description **1**
1.1 Microgrid Controller 1
1.2 OpenADR Standard 1
2 High-Level Sequence of Operation **3**
2.1 Estimation of the Desired Load-Shed Requirement 4
2.2 Additional Considerations 6
3 Communication Specifications **8**
3.1 OAD Power Threshold Setpoint: Method 1 8
3.2 MGC/OAD Interactive Messaging: Method 2 10
3.3 Simple Demand Reduction Request: Method 3 13
3.4 Granular On/Off Control: Method 4 15
4 Required Changes to Microgrid Controllers **18**
4.1 Microgrid Controller Modifications 18
4.2 Microgrid Simulation Environment Modifications 19
5 Preliminary Test Plan **20**
5.1 Grid Load Increases and Approaches Grid Spinning Capacity: Test Case 1 20
5.2 Grid Source Becomes Unavailable (Spinning Capacity Reduced): Test Case 2 21
5.3 Grid Spinning Capacity Increased While Loads Have Been Shed by OAD: Test Case 3 21
5.4 Closed Transition to Island with Insufficient Islanded Spinning Capacity: Test Case 4 21
5.5 Transition to Grid-Tied While Loads Have Been Shed by OpenADR: Test Case 5 22
5.6 Black Start to Islanded Mode with Insufficient Reliable Sources: Test Case 6 22
References **23**

List of Figures

Figure 1. Data flow in proposed method ... 3
Figure 2. Interaction between MGC and OAD for OAD power threshold setpoint ... 10
Figure 3. Interaction between MGC and OAD for MGC/OAD interactive messaging ... 11
Figure 4. Interaction between MGC and OAD for granular on/off control ... 16
Figure 5. Example grid topology for testing and validation ... 20

List of Tables

Table 1. OAD with kW Threshold Message Contents ... 9
Table 2. Proposed Modbus Registers for OAD Power Threshold Control Data Exchange ... 9
Table 3. Request for Demand Reduction Message Contents ... 12
Table 4. Verification of Receipt of Demand Reduction Request Message Contents ... 12
Table 5. Request Cancellation Message Contents ... 13
Table 6. Confirmation/Denial That Request Can Be Achieved Message Contents ... 13
Table 7. Request for Demand Reduction Message Contents ... 14
Table 8. Proposed Modbus Registers for MGC/OAD Data Exchange ... 14
Table 9. Message Contents for Individual Load Management ... 15
Table 10. Proposed Modbus Registers for Granular On/Off Control Data Exchange ... 17

1 Project Description

When military microgrids isolate themselves from the main electrical grid, they must locally balance electricity supply and demand. 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 considered noncritical may lose service entirely. The remotely controlled switchgear needed to manage load in this way is expensive ($30,000 to $50,000 per building, installed). While effective at shedding load, this strategy disrupts installation operation and risks damaging equipment during both disconnection and re-energization.

With the goals of lowering costs, protecting equipment, and enhancing the agility of U.S. Department of Defense microgrids, this project demonstrates the use of cybersecure automated demand response (ADR) technology to manage microgrid loads during islanded mode of operation. This automated approach achieves load shedding and shifting through communication signals sent to equipment controllers rather than by cutting off the flow of electricity within the microgrid itself. 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.

1.1 Microgrid Controller

This project utilizes the GridMaster® microgrid controller (MGC) from S&C Electric Company for this demonstration. GridMaster is a combined software and hardware platform built for advanced microgrids and is agnostic to the connected equipment, facilitating integration and communication with a host of different distributed energy assets. It uses grid state data to determine the appropriate actions to take and then commands the microgrid assets during both grid-tied and islanded modes. GridMaster is based on a distributed architecture and, consequently, it provides inherent redundancy for more reliable operation. At any given time, one GridMaster controller serves as the lead decision maker by requesting data from the other GridMaster controllers and making control decisions based on the observed state of the grid (S&C Electric Company). The developed methods are generic enough to be adopted by other microgrid controllers.

1.2 OpenADR Standard

This project employs automated demand response (ADR) technology to facilitate communication between a microgrid controller and a BMS. Many ADR solutions are built upon the OpenADR communications protocol, which originated in 2002 (OpenADR Alliance, 2015). To support this development, the OpenADR Alliance was established in 2010 by a consortium of utilities, software vendors, device manufacturers, national labs, DR aggregators, testing and certification labs, system integrators, and consulting firms (OpenADR Alliance)The alliance currently boasts over 200 members and has been instrumental in the creation and adoption of OpenADR, with over 280 commercial products now compliant with its specifications (OpenADR Alliance) . This research aims at expanding existing communication technologies for automating demand response in microgrid applications.

The purpose of this document is to outline the protocols and data exchange for MGC interaction with an OpenADR appliance.

2 High-Level Sequence of Operation

The MGC coordinates the operation of microgrid components, including power sources, switchgear, energy storage, and loads. There are two primary goals of the controller when the microgrid is islanded: (1) ensure that critical loads remain energized, and (2) minimize fuel consumption to prolong the ability of the grid to continue operating. These two goals are antithetical to each other in that the best way to ensure that all critical loads remain energized is to operate as many power sources as possible—which, of course, will consume the most fuel.

In practice, the MGC compromises between the two goals. The user can define priorities on the various loads such that lower-priority loads will be shed as the overall grid demand approaches the available generation capacity. Traditionally, MGCs have shed the lower-priority loads by either communicating directly with a load control device or by operating switchgear to de-energize entire sections of the grid. It is recognized that these are coarse methods of load management, and it is believed that using OpenADR will provide more acceptable and rational methods of shedding load.

This section describes a proposed high-level sequence of operation that combines an MGC with an OpenADR appliance to provide a more efficient approach to load reduction in an islanded microgrid. The overall system data flow is shown in Figure 1.

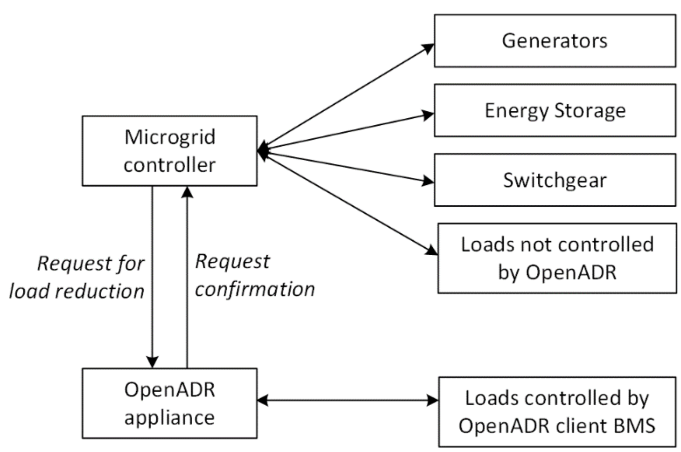

Figure 1. Data flow in proposed method

In microgrids, load shedding occurs when one of two conditions are met. Loads may be shed when electrical demand approaches the generation capacity of all usable power sources. This is a last resort of the system; the system will always attempt to bring on more sources rather than to shed load. The second condition is when a power source becomes unusable due to a fault, fuel exhaustion, required maintenance, or some other reason. If a reliable source becomes unusable (as detected by an error condition or if marked unavailable by the user), and there is no other source to bring online, loads may be shed to safely curtail the affected power source.

2.1 Estimation of the Desired Load-Shed Requirement

The MGC determines when load shedding is necessary in the microgrid. The amount of load to be shed is calculated from the spinning reserve available on the grid and includes various safety ratios and buffers. In this context, “spinning reserve” is the remaining capacity of reliable sources that are operating but only partially loaded. For example (and ignoring safety buffers and capacity ratios), a 100-kW generator producing 40 kW has 60 kW of spinning reserve. In practice, the reserve calculations consider both capacity ratios and safety buffers. These reserve calculations set a threshold for both power source dispatch/curtailment and load shedding/restoration. For power sources, the capacity ratios associated with power source dispatch will be higher than that for the power source curtailment. (And the safety buffers will be smaller for power source dispatch than that for power source curtailment).

For example, a source dispatch capacity ratio of 0.8 implies that a power source will dispatch once the electrical demand rises above 80% of the available generation capacity. This would be coupled with a source curtail capacity ratio that is smaller than the dispatch capacity ratio (e.g., 0.75) to introduce hysteresis. In other words, new power generation sources will be dispatched as the demand rises above a given threshold, but the sources will not be curtailed again until the demand drops below a different threshold that is lower than the dispatch threshold.

Similarly, there are different capacity ratios used for load shedding and load dispatch. Note that load shedding is only done if there is no additional generation that can be dispatched. Loads are shed if total demand rises above a given threshold, but they are not restored until demand drops below a lower threshold.

An important nuance is that the capacity ratios reference the ratio of the electric load to the capacity of reliable sources. Reliable in this context refers to nonrenewable power sources such as conventional (fossil-fueled) generation sources and energy storage systems. Renewable sources are considered too intermittent and unreliable to be included in the capacity ratio calculation.

The reserve of the system is calculated as the smaller of (1) the difference between the total capacity and the demand plus a safety buffer, and (2) the difference between the total capacity corrected by the capacity ratio and the demand.

$$Reserve = min\left[C - \ (L + B), R_{cap} \ x \ C - L\right]$$

where

C = the generating capacity of the system
L = the load on the system
B = a safety buffer, and
R_{cap} = the capacity ratio

For instance, given a capacity ratio of 0.9 and a buffer of 10 kW, the resulting reserve of a system with a single 100-kW generator running at 40 kW would be:

$$Reserve = min[100 - \ (40 + 10), 0.9\ x\ 100 - 40] = 50\ kW$$

Note that buffers are used for the entire system and not on a source-by-source basis.

Both short-term and long-term reserves are considered when determining desired load-shed magnitudes. The distinction is made using different ratios and buffers. It may be acceptable to go over the rated capacity of a generator for a short period of time. As such, if the load should increase and push the generator past its rated capacity but not past its short-term capacity, then there is a window to bring another source online or perhaps shed load. The overall reserve is the minimum of the long-term reserve or the short-term reserve:

$$Overall\ \ Reserve = min(long\ term\ reserve, short\ term\ reserve)$$

For instance, if a single source, which is rated at 100kW and output 120kW for a short-term, is currently outputting 40 kW; and has the following given safety ratios and buffers:

long-term capacity ratio: 0.9
long-term buffer: 10 kW
short-term capacity ratio: 0.8
short-term buffer: 20 kW

The minimum reserve would be:

$$Reserve = min[100 - (40 + 10), 0.9\ x\ 100 - 40, 120 - (40 + 20), 0.8\ x\ 120 - 40] = 50kW$$

As stated previously, reserve calculations come in pairs. There is a source dispatch reserve calculation, which is paired with a source curtailment calculation. By having more conservative values for curtailment than for dispatch, a hysteresis band results. For example, a source may be dispatched when the load exceeds 90% of capacity but only curtails if the load falls below 80% of capacity.

Reliable sources that do not load-follow and instead need to have their output set (e.g., batteries or generators in baseload mode) cannot utilize their full capacity over the short term. If the loads suddenly increase, these sources will not adjust their output in a timely fashion and instead will allow other sources to be overloaded. Thus, for short-term reserve calculations, these sources can only rely on their current output. The rated capacity can be used in long-term reserve calculations, as the system will adjust these sources as loads change.

When the reserve calculation is applied to load shedding, the amount of load to be shed is the calculated reserve amount. If the reserve calculation result is negative, then this number represents the minimum magnitude of load that needs to be shed. If the reserve calculation result is positive, then this number represents the magnitude of load that can be restored.

There are two additional attributes assigned to loads that influence shedding and restore behavior. The first is the priority of the load. This is a user-defined value that indicates the importance of a given load relative to the other loads in the grid. Lower-priority loads are shed before higher-priority loads, even if multiple lower-priority loads must be shed to match the demand of a single higher-priority load. When restoring loads, higher-priority loads are brought back online before the lower-priority loads. The only exception to this rule is when there are not enough lower-priority loads to meet the demand reduction target and it might be necessary to shed a higher-priority load to ensure that the micrgrid does not collapse.

The second additional attribute is the anticipated demand assigned to each load. This is a user-defined value that is used when restoring loads that have been shed. It is not sufficient to assume the kW of a given load will be the same before and after being shed. Indeed, for many loads such as HVAC and industrial processes, the demand during initial start-up can be significantly higher than that during steady-state operation. The user-entered anticipated demand value represents the highest reasonable kW expected once the load is restored. This anticipated demand is compared to the calculated reserve values when determining if a given load can be restored.

2.2 Additional Considerations

When operating in droop, the reserves are determined by following the droop curves of the reliable sources online and calculating the maximum load those sources can handle. Droop is a control strategy to regulate the output of parallelly connected generators in a power generation system. For example, for short-term calculations, the droop curves of the reliable online sources are followed until any source exceeds its short-term capacity. The load that the sources could handle prior to overloading any source is the 'capacity' of that group of equipment (at those offsets). Note that this is different from simply adding up the rated capacities of the sources, since different power sources may be at different spots in their droop curves (based on frequency offsets), and the system is consequently limited in the power it can produce at the given grid frequency. For instance, generators are often adjusted to get them running at their minimum values, while at the same time, the frequency offset of energy storage systems may be adjusted to force them to charge. In this scenario, the generators are more prone to reach their maximum output (and are closer to being overloaded) than the energy storage.

The microgrid can also be configured to operate additional capacity for the purposes of energy resiliency. This is often the case with Department of Defense installations. In so-called "N" mode, only enough generators are dispatched to meet the loads. While this saves fuel, it also means that the loss of any one source could result in grid voltage collapse. In "N" mode, reserve values are calculated using all available reliable spinning sources. Also, in "N" mode, it is assumed that any renewable sources will trip off with an unanticipated stoppage of the first source, and consequently, the reserves do not include any renewable capacity.

In "N+1" mode, the MGC attempts to keep sufficient reliable sources online such that any single power source can stop, and the remaining sources will have sufficient capacity to prevent an outage. This criterion is a short-term consideration, since eventually another source can be dispatched and/or loads can be curtailed. Therefore, multiple short-term reserve values are calculated, each with a different source removed from the spinning sources. The reserve is the minimum value produced from these permutations and from the "N" based long-term calculation. This strategy requires that at least two reliable sources are running when in island mode. As with "N" mode, renewables are assumed to trip off with an unanticipated stoppage of the first source, and so the reserve calculations do not include any renewable capacity.

In "N+2" mode, the MGC attempts to keep enough reliable sources online such that should any two sources stop, the remainder can satisfy the grid demand. The behavior and calculations in "N+2" mode are similar to that of "N+1" mode, except this mode requires at least three reliable sources to be running when in island mode.

3 Communication Specifications

This section describes the necessary communication between the MGC and the OpenADR device (OAD). 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 pre-established strategies or user settings. They also provide feedback to the VTN, allowing for the monitoring and coordination of the demand response event.

It is important to recognize that the MGC itself does not use the OpenADR protocol but rather will be communicating desired demand reduction values to the OAD. The OAD, in turn, will exchange compliant OpenADR messages with VENs to implement the requested demand adjustment. In this project, Modbus is used as a communication protocol between the MGC and the OAD. Adaptation to other protocols should be relatively straightforward.

While it is expected that communications between the VTNs and VENs comply with the OpenADR protocol, there could be variations in how the VTNs communicate with higher-level devices. A few different options are presented below that can be equally applied to different MGCs depending on their capabilities. The first method discussed is preferred, in part because it does not expose information about the loads. The second method is based on a more direct MGC/OAD communications and assumes an API-like interaction between them. The third and fourth methods focus on Modbus-centric solutions that provide somewhat similar functionality to the first. Note that for the methods below that use Modbus, there would need to be some form of translation between the Modbus messaging and the OAD so that the messages between the MGC and the OAD follow the appropriate message format.

3.1 OAD Power Threshold Setpoint: Method 1

This method can be used if the OAD has the ability to monitor the loads under its control and manage those loads so that they do not exceed a threshold established by the MGC. The proposed data fields are described in Table 1. In this method, the MGC writes to a register that indicates the maximum allowable kW for the loads. This value is updated continuously by the MGC and is a function of the grid conditions.

The OAD reads this maximum allowable value and then manages the loads to ensure that the threshold is not exceeded. It will shed loads if necessary and will restore loads if there is available and sufficient capacity to do so.

The OAD will provide a *Baseline load* value, which indicates the load the system would use if it was not in an event. This would be used by the MGC to potentially dispatch additional reliable sources to allow those loads to be dispatched. (This would be communicated from the MGC to the OAD through an increased *Allowed load threshold* value.)

Table 1. OAD with kW Threshold Message Contents

Field	Units	Data Type	Description
Allowed load threshold	kW	Float	Maximum allowable kW as determined by the MGC
Load hysteresis	kW	Float	Difference between the allowed load threshold when loads are curtailed and the threshold when loads are restored, in kW, written by the MGC
Load power	kW	Float	Written by the OAD indicating the present value of the load as measured by the OAD
Available load shed	kW	Float	An optional value written by the OAD indicating the amount of load that can be shed
Baseline load	kW	Float	Amount the OpenADR system would use if it was not in an event

This method assumes that the OAD has active load management capabilities and can respond to the threshold set by the MGC. It may also be desirable to include two values for the allowable kW threshold: one for when the loads are increasing and one for when they are decreasing. The two values would create a hysteresis that could help prevent undue curtailment/dispatch of loads if the electric demand was hovering around a single threshold. This is accounted for in the proposed Modbus map in Table 2 and illustrated in the flow chart of Figure 2.

This method does not require the OAD to have visibility into all the loads on a given circuit, since presumably the MGC can perform the math necessary to calculate the portion of the total load that is managed by the OAD. It does, however, require an upfront commissioning effort on the MGC to establish that distinction.

Table 2. Proposed Modbus Registers for OAD Power Threshold Control Data Exchange

Reg	Name	Description
40001	Upper allowed load threshold	Max load to be allowed by OAD in kW; above this value, loads need to be shed; written by MGC
40002	Hysteresis margin	(optional) A value in kW that represents the margin between the value in register 40001 and the point at which loads are restored; written by MGC
40003	Load power	Present value of load in kW as measured by the OAD; written by OAD
40004	Available load shed	(optional) The estimated amount of load that can be shed as determined by the OAD, in kW; written by OAD
40005	Baseline load	Amount the OpenADR system would use if it was not in an event; written by OAD

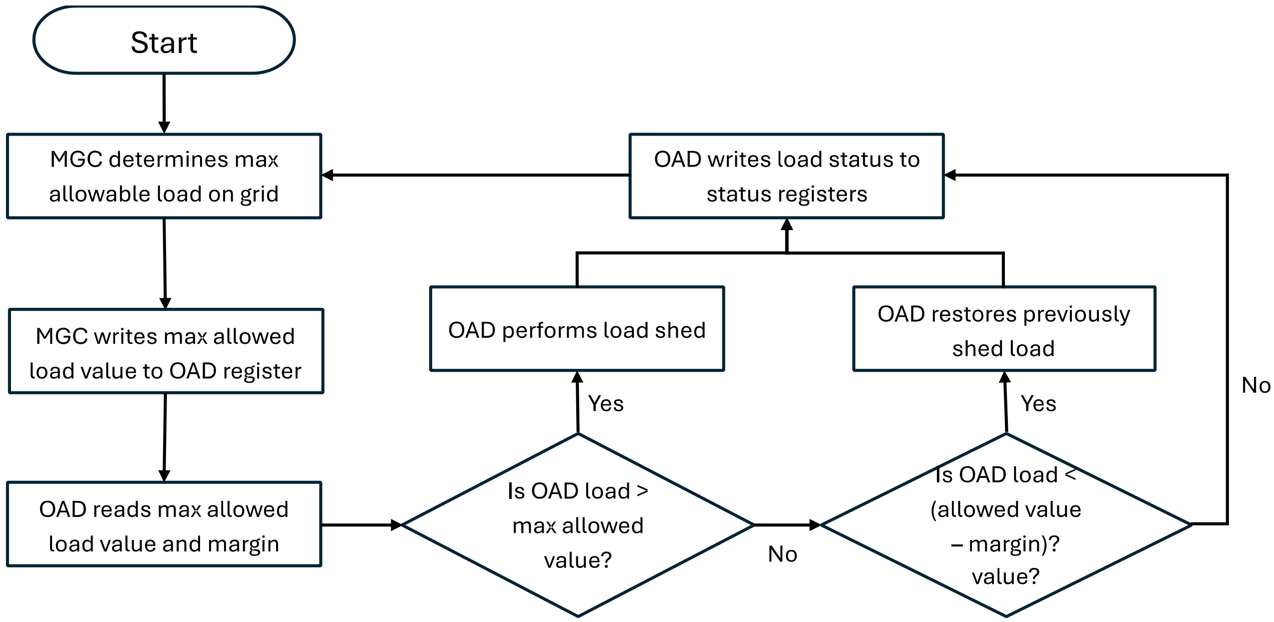

Figure 2. Interaction between MGC and OAD for OAD power threshold setpoint

3.2 MGC/OAD Interactive Messaging: Method 2

This method assumes an interaction between the MGC and the OAD, as summarized in the flow chart in Figure 3. This method is meant for interactive messages where the format would ultimately be based on protocols that are supported by the MGC and OAD.

Once a target load-shed value is identified by the MGC, the value is sent to the OAD. To prevent multiple, repeated messages, the OAD confirms it has received the request and provides an indication of the status of the request. The MGC waits for a response and then attempts to resend the request. If there is no response from the OAD after a certain number of tries (or after a timeout period has expired), then the MGC broadcasts a cancel message to the OAD and assumes that the original request will not be satisfied. The MGC then attempts to perform load shedding through a more traditional method such as opening entire building switches, if available.

Once the OAD successfully confirms receipt of the load shed request, the MGC starts a timer to determine if the load shed occurs within an acceptable time window. At the same time, the OAD determines if the load-shed request can be honored given the current state of the system and any occupant preferences. If the OAD determines that the request cannot be honored, it notifies the MGC, which then executes load shedding through more traditional methods such as opening one or more switches to disconnect enough load from the grid.

If the OAD acknowledges the request and indicates it is working on satisfying the request, the MGC tracks the measured load for the expected reduction. If the load does not decrease within the allowed time window, the MGC implements load shedding through traditional methods.

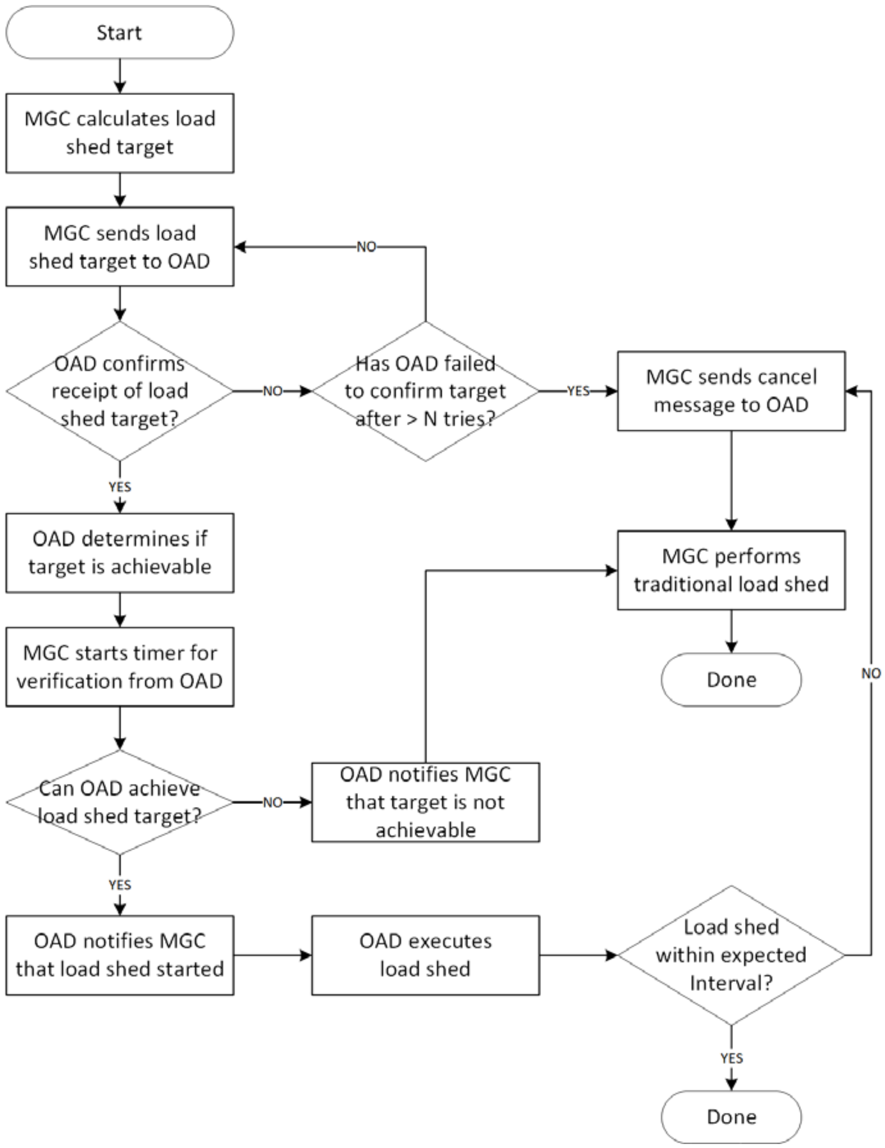

Figure 3. Interaction between MGC and OAD for MGC/OAD interactive messaging

When sending a request for demand reduction, the MGC uses the fields described in Table 3. It is recognized that this message may not be necessary, since canceling an event that is not acknowledged in the first place may not initiate any action by the OAD side. However, it is included here in case the OAD has received the original request, and the lack of acknowledgement is due to an issue with communication. These values are sent from the MGC to the OAD. The *Message ID* is used here (and in the other tables) to uniquely identify a given request or response. It is expected that the ID will be used by the OAD to prevent shedding too much load if the same message is repeated. The *Target demand reduction* is the value of the requested load shed in kW. This value is calculated by the MGC based on the reserve calculations described earlier. The *Time of request* is the timestamp of when the MGC created and broadcast the message to the OAD and acts as time zero for the expected time window over which the load shed must occur. Since the MGC is calculating desired load shed based on the reserve calculations, there is a safety margin during which the load shed can occur, typically a few minutes. Since the load-shed request is usually triggered by increasing demand, there is the likelihood that the demand will continue to rise and could potentially overwhelm the existing generation if the load shedding does not occur within some window. The size of this window is also transmitted to the OAD in this message. If the OAD is unable to achieve the target load shed within this time frame, it should cancel the attempt, since it is assumed that the MGC will take other actions to drop load to ensure that the grid remains energized.

Table 3. Request for Demand Reduction Message Contents

Field	Units	Data Type	Description
Message ID	<n/a>	Long	Unique identifier that can be used to avoid duplicate processing of repeated messages
Target demand reduction	kW	Float	Value of required demand reduction
Time of request	Time/Date	Long	Linux time representation of when request was sent by MGC
Allowable window for request handling	Seconds	Integer	The demand reduction must be achieved within this time

Once the OAD receives the request for demand reduction described in Table 3, it needs to confirm receipt. Otherwise, the MGC may attempt to do more traditional load shedding and the grid could experience more loss of load than necessary. The fields in Table 4 are used for this confirmation. These values are sent from the OAD to the MGC. This message simply reflects the values in the original request, i.e., the *Request message ID* and the *Target demand reduction* should be the same as those in the requesting message.

Table 4. Verification of Receipt of Demand Reduction Request Message Contents

Field	Units	Data Type	Description
Message ID	<n/a>	Long	Unique identifier that can be used to avoid duplicate processing of repeated messages
Request message ID	<n/a>	Long	Message identifier of original request from MGC
Target demand reduction	kW	Float	Confirmation of requested reduction

If the OAD does not confirm receipt of the demand reduction request (or if the confirmation is not received within a given time-out period), the MGC will attempt to cancel the original request. This is done using the fields shown in Table 5, which are sent from the MGC to the OAD. Given that the lack of confirmation from the OAD may represent a communication issue or a hardware problem, the MGC will not expect a response to the cancellation message. In essence, this is a "fire-and-forget" message, and once it is sent, the MGC may attempt to perform load shedding through more traditional methods.

Table 5. Request Cancellation Message Contents

Field	Units	Data Type	Description
Message ID	<n/a>	Long	Unique identifier that can be used to avoid duplicate processing of repeated messages
ID of original request	<n/a>	Long	Identifier of message with original demand reduction request that is to be cancelled

There may be occasions when the OAD cannot satisfy the demand reduction request from the MGC. In this case, the OAD will send the information in Table 6 back to the MGC. This table contains information similar to that of the previous messages with the addition of the *Achievable* field, which is set to true if the request can be accomplished and false otherwise. This table also has an optional *Alternative target* field that may be set by the OAD to whatever potential demand reduction is possible.

Table 6. Confirmation/Denial That Request Can Be Achieved Message Contents

Field	Units	Data Type	Description
Message ID	<n/a>	Long	Unique identifier that can be used to avoid duplicate processing of repeated messages
Request message ID	<n/a>	Long	Message identifier of original request from MGC
Target demand reduction	kW	Float	Confirmation of requested reduction
Achievable	True/False	Boolean	Verification that request can be honored
Alternative target *(optional)*	kW	Float	Estimate of achievable load reduction at the present time

3.3 Simple Demand Reduction Request: Method 3

This proposed method represents the simplest communication technique between the MGC and the OAD, with the contents of messages between the two summarized in Table 7. In this case, the MGC indicates that the OAD should shed load by (1) setting the *Demand reduction requested* point to TRUE and (2) setting the required demand reduction amount in the *Target demand reduction* point.

Once the OAD has shed the requested amount of load, or when it has shed all the load it can (if it can't meet the requested amount), the OAD sets the *Demand reduction requested* field to FALSE. The MGC consistently monitors the *Demand reduction requested* value and recognizes that the load-shed process has completed when it changes to FALSE. Note that this strategy requires that the *Demand reduction requested* value is stored in a Read/Write register with both sides able to modify the value at appropriate times. In comparison, the *Target demand reduction* point is only ever written by the MGC and read by the OAD. Load dispatch follows a similar process, but with the analogous points *Demand increase requested* and *Target demand increase allowed*.

Table 7. Request for Demand Reduction Message Contents

Field	Units	Data Type	Description
Demand reduction requested	True/False	Boolean	True/False whether load shed is currently requested or not
Target demand reduction	kW	Float	Value of required demand reduction
Demand increase requested	True/False	Boolean	True/False whether load dispatch is currently requested or not
Target demand increase allowed	kW	Float	Requested load increase allowed in kW

A potential representation of the Modbus maps corresponding to this method is given in Table 8. The values shown here are listed as holding registers, although the True/False values could be coil registers as well.

Table 8. Proposed Modbus Registers for MGC/OAD Data Exchange

Reg	Name	Description
40001	Demand reduction requested	True/False whether load shed is currently requested or not 0/False: no request 1/True: request for load shed as specified in register 40002
40002	Requested demand reduction	Requested load shed, kW
40003	Demand increase requested	True/False whether load can be added 0/False: no request 1/True: request for load shed as specified in register 40004
40004	Target demand increase allowed	Load increase allowed, kW

Note that the values in Table 8 are the bare minimum necessary for this method to be applied. There are several other useful data fields that could be written by the OAD, but these are contingent on the ability of the OAD to determine and report on these values. They include:

- Maximum Sheddable Load: This is the maximum amount of load that could be shed by the OAD, regardless of the request from the MGC.

- OAD Status: This reports the operating condition of the OAD and consists of a bit field that enumerates the error status of the OAD, whether it is active and able to shed loads, etc.

- Load Amount: This would be the reported total load from the OAD that would be used to compare what the OAD believes the load is versus what the MGC believes the load is.

- Requested Load: This value represents the total shed load that could be restored, assuming there was sufficient available generation capacity in the microgrid.

3.4 Granular On/Off Control: Method 4

In this method, the MGC is closely coordinated with the OAD and is "aware" of all the individual loads that can be managed by the OAD. For each individual load, there will be a series of control points (see Table 9) commanded by the MGC and subsequently implemented by the OAD. This method assumes on/off control of each load and would have to be modified if the loads were to be controlled proportionally.

The data streams associated with each load are as follows. The interaction between the MGC and OAD is illustrated in Figure 4.

- The load kW is measured by the OAD, and the communication register is continuously updated with the present value of the load. This value is read by the MGC and can be used to determine if the load should be shed.
- An on/off command is written by the MGC when there is a need to shed a particular load. This value is then read by the OAD which, in turn, performs the actual load shed/restore.
- The on/off status is written by the OAD and indicates if the request to shed that load has been satisfied. It will be set back to the default (on) value by the OAD when the load is restored.
- An error status is an optional field written by the OAD that provides an explanation if the on/off command cannot be honored. Reasons for this might be user override, inability to communicate with the load, etc.

If using this method, the MGC will require a commissioning effort where each controllable load is assigned an anticipated load and a load priority. The anticipated load is the maximum expected value of the load that would occur immediately after restoration. In the case of any HVAC loads, this would include the initial pull-down power draws that occur when compressors and other similar equipment first start up.

The load priorities are used-set values that assign the relative importance of each load to the others. The priorities can be in simple categories (e.g., low, medium, and high) or assigned individual ranks. Note that load priorities can and will change over time depending on facility mission and use on any given day, so the priorities cannot be considered static "read once" values.

Table 9. Message Contents for Individual Load Management

Field	Units	Data Type	Description
Load 1 power	kW	Float	The power consumption of load 1 as measured by the OAD
Load 1 on/off command	True/False	Integer	Command written by MGC indicating if load 1 should be shed
Load 1 on/off status	True/False	Integer	Status written by OAD on whether load 1 has been shed

Field	Units	Data Type	Description
Load 1 error status	True/False	Integer	An optional field that is an enumerated value of potential errors related to the load shed attempt
Load 2 power	kW	Float	The power consumption of load 2 as measured by the OAD
Load 2 on/off command	True/False	Integer	Command written by MGC indicating if load 2 should be shed
Load 2 on/off status	True/False	Integer	Status written by OAD on whether load 2 has been shed
Load 2 error status	True/False	Integer	An optional field that is an enumerated value of potential errors related to the load shed attempt
Pattern is repeated for as many loads controlled by OAD			

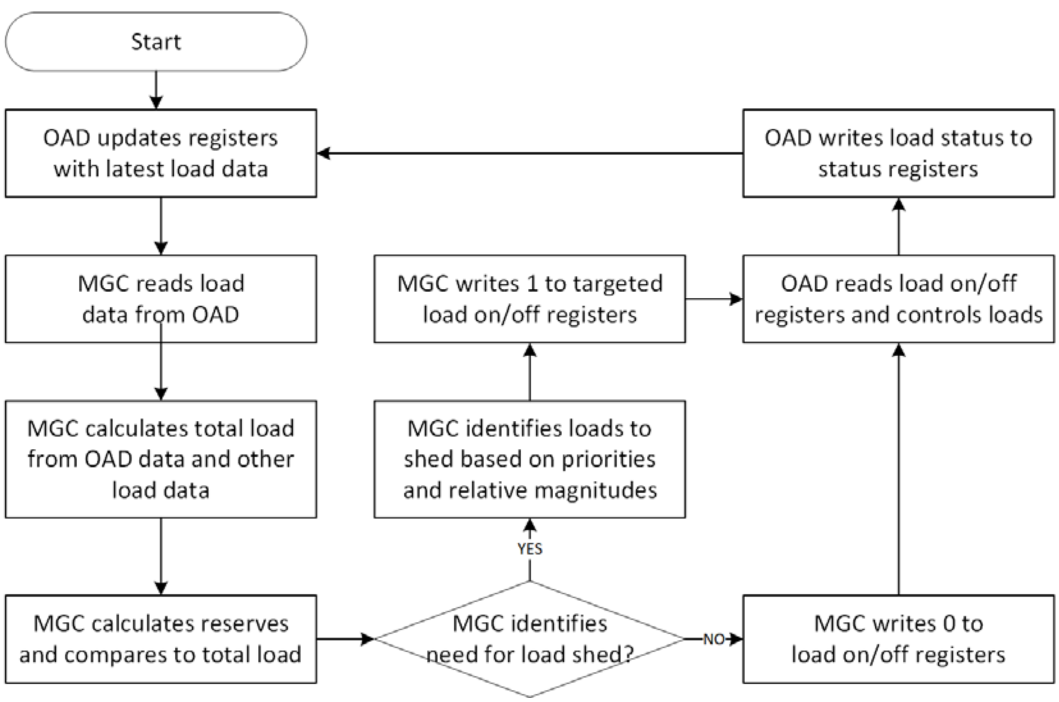

Figure 4. Interaction between MGC and OAD for granular on/off control

The proposed Modbus map for this method is given in Table 10.

Table 10. Proposed Modbus Registers for Granular On/Off Control Data Exchange

Reg	Name	Description
40001	Load 1 power	Power consumption of load 1 in kW
40002	Load 1 on/off command	Command to turn load 1 on/off
40003	Load 1 on/off status	On/off status of load 1
40004	Load 1 error status	(optional) Flag indicating error
Pattern is repeated for all subsequent loads		

4 Required Changes to Microgrid Controllers

The topologies and control capabilities of many microgrids dictate that demand reduction is achieved through traditional methods, specifically operating switchgear that controls power flow at the building level or at the individual load level (if there is supporting hardware to do so). With the methodology proposed in this report, the MGCs can first attempt to accomplish load shedding through communication with the OAD. If the communication with the OAD is not successful for any reason, the controller can then operate the appropriate switch gear under its direct control.

4.1 Microgrid Controller Modifications

This project utilizes commercial server and client ADR products from GridFabric (GridFabric, n.d.), and the required modifications to an MGC will need to consider the communication methods supported by the associated OpenADR appliance, typically a VTN. The first proposed communication method may require adding support for the APIs from the VTN. Of the remaining three methods, it may be necessary to adapt any existing load shed/add calculations in the existing controller. The functionality required for all four proposed methods can likely be implemented through the user-defined processes. The method changes listed below include suggestions for new custom processes that the user can write.

- Method 1: As mentioned above, this would require wholesale changes to the controller communication libraries to add support for the OpenADR appliance API.

- Method 2: For this method, the controller changes include adapting any existing load-shed/add algorithms and creating custom processes to (1) calculate the maximum allowable total OAD load; (2) write the calculated maximum kW value to the OAD; (3) integrate the *Baseline load* value into the existing *Total requested load* calculations to dispatch sources to bring on OAD loads; (4) monitor the OAD assets' power level; (5) implement a OAD power-exceeded timer; and (6) add routines to issue load shedding commands to other loads not controlled by the OAD in the event that the OAD asset power levels exceed the issued maximum allowable total OAD load for some user-definable duration.

- Method 3: This method would necessitate removal of any default load-shed/restore algorithms, and the creation of custom processes to (1) calculate the desired load-shed/restore kW, (2) issue load-shed/restore requests to the OAD device connection, and (3) implement an OAD load-shed/restore request timer. If the OAD load-shed/restore requests time out or communications with the OAD are lost, then the controller would also need to remove/reset the OAD load-shed/add request(s) prior to shedding other loads not controlled by the OAD.

- Method 4: In this method, the OAD load assets can be treated like the existing load types in the controller.

4.2 Microgrid Simulation Environment Modifications

Simulation programs are often used to test communication paths and for vetting new and existing control algorithms. This testing is crucial for ensuring that the system is behaving correctly prior to installation and commissioning of the equipment and controller. For testing OpenADR solutions, it may be necessary to modify existing simulations to accommodate OpenADR messages and the load response to strategic load shed requests.

5 Preliminary Test Plan

A reference test grid such as that shown in Figure 5 needs to be created that represents a typical Department of Defense installation. It should include a mix of load priorities as well as both conventional and renewable power sources. The power sources should be sized, and daily load profiles should be chosen such that the peak demand cannot be satisfied even when all power sources are dispatched. This will require that some loads be shed. The proposed test scenarios are described below.

5.1 Grid Load Increases and Approaches Grid Spinning Capacity: Test Case 1

Description: On an islanded grid, if the grid load increases and/or the island spinning reserve decreases to a predetermined threshold, the MGC shall send a command to the OAD to shed some load. This test case includes two separate scenarios/tests:

1. The OAD sheds the requested kW
2. The OAD is unable to shed the requested kW

Successful Outcome: This test is considered successful if any of the following happens:

1. The OAD sheds the requested kW, the microgrid spinning reserve rises above the predetermined threshold, and no additional loads are required to be shed.
2. The OAD sheds some load but does not shed the requested kW. The MGC sheds other loads and the microgrid spinning reserve rises above the predetermined threshold, and no additional loads are required to be shed.

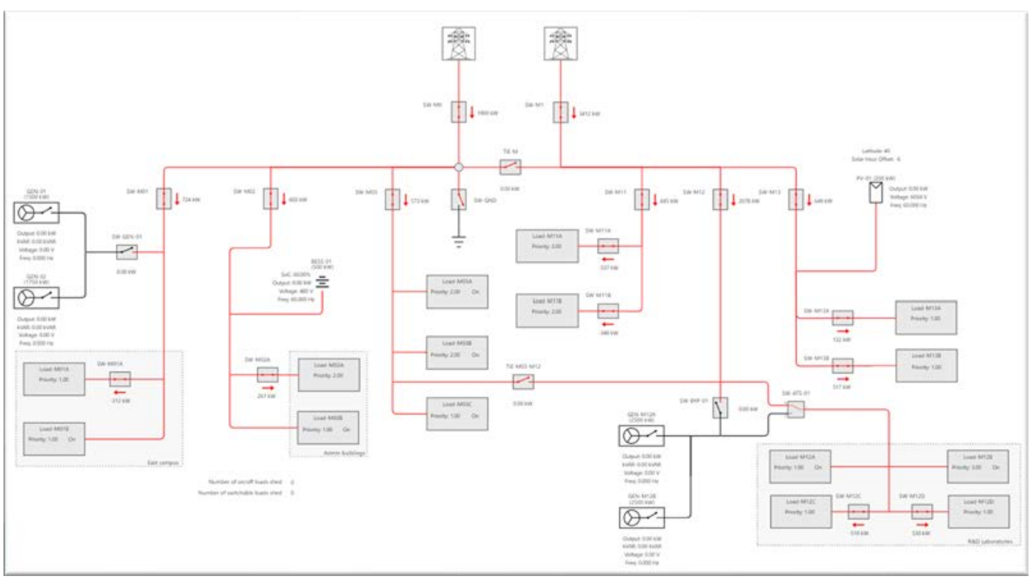

Figure 5. Example grid topology for testing and validation

5.2 Grid Source Becomes Unavailable (Spinning Capacity Reduced): Test Case 2

Description: On an islanded grid, if an online reliable source (battery energy storage system or generator) is marked unavailable, the microgrid spinning reserve is effectively reduced. If this reduction in spinning reserve goes below a predetermined threshold, the MGC sends a command to the OAD to shed some load. This test case includes two separate scenarios/tests:

1. The OAD sheds the requested kW
2. The OAD is unable to shed the requested kW

Successful Outcome: This test is considered successful if any of the following happens:

1. The OAD sheds the requested kW, the microgrid spinning reserve rises above the predetermined threshold, and no additional loads are required to be shed.
2. The OAD sheds some load but does not shed the requested kW. The MGC sheds other loads and the microgrid spinning reserve rises above the predetermined threshold, and no additional loads are required to be shed.

5.3 Grid Spinning Capacity Increased While Loads Have Been Shed by OAD: Test Case 3

Description: On an islanded grid with loads shed by OAD, if the microgrid spinning capacity increases enough to support additional loads, the MGC shall send a command to the OAD to add loads.

Successful Outcome: This test is considered successful if the OAD adds loads and the microgrid spinning capacity remains above the determined threshold for load shedding.

5.4 Closed Transition to Island with Insufficient Islanded Spinning Capacity: Test Case 4

Description: During a closed transition from grid-tied to islanded, if the microgrid spinning reserve is below the determined threshold for load shedding, the MGC shall send a command to the OAD to shed loads ***prior to*** opening the utility breaker. This test ***does not*** consider the case where loads need to be shed in addition to the loads shed by the OAD.

Successful Outcome: This test is considered successful if the following hold true:

1. The island state is reached
2. After islanding:
    a. The MGC does not need to immediately shed additional loads. This assumes a reasonably constant load immediately before, during, and immediately after the closed transition.

b. The microgrid spinning reserve is above the determined threshold for load shedding.
c. The OAD has shed the minimum possible amount of kW to keep the microgrid spinning reserve above the determined threshold for load shedding.

5.5 Transition to Grid-Tied While Loads Have Been Shed by OpenADR: Test Case 5

Description: After a closed transition from islanded to grid-tied, the OAD will add back all loads that have been shed while islanded.

Successful Outcome: This test is considered successful if the following hold true:

1. The grid-tied state is reached
2. After becoming grid-tied (and not before then), the OAD has added back all loads under its control

5.6 Black Start to Islanded Mode with Insufficient Reliable Sources: Test Case 6

Description: During a black start, the MGC will determine if the full load of the OAD assets can be added to the microgrid and will send the appropriate load shed commands to the OAD (0 kW to some positive kW value) ***prior to*** energizing the OAD circuit.

Successful Outcome: This test is considered successful if the following hold true:

1. Black start is successfully completed
2. Once islanding has completed and assuming a relatively steady grid load, the MGC will not need to immediately shed any loads.
3. Assuming a total microgrid load that exceeds the microgrid spinning reserve, the OAD will have shed loads prior to the OAD circuit being energized. Note that this isn't a true load shed because, in this case, the OAD assets are not energized yet. Rather, the OAD is instructed to open certain load circuits or somehow ensure that those loads are not energized when the OAD circuit is energized.

References

GridFabric. “GridFabric.” Accessed July 16, 2024. https://www.gridfabric.io.

Modbus “Modicon Modbus Protocol Reference Guide” , https://www.modbus.org/docs/PI_MBUS_300.pdf.

OpenADR Alliance. 2015. "OpenADR 2.0 Profile Specification B Profile."

OpenADR Alliance. “OpenADR Alliance.” Accessed July 16, 2024. https://www.openadr.org.

S&C Electric Company. “GridMaster Microgrid Control System.” Accessed July 16, 2024. https://www.sandc.com/en/products--services/products/gridmaster-microgrid-control-system/.