Vulcan Test Platform: Demonstrating the Data Center as a Flexible Grid Asset

Deepthi Vaidhynathan,[1] Kumaraguru Prabakar,[1] Gregory Martin,[1] Anand Ramesh,[2] Ben Wheeler,[2] Christopher Coco,[2] Jimmy Clidaras,[2] Matthieu Monsch,[2] Sangsun Kim,[2] Saurav Talukdar,[2] and Shilpa Marti[2]

1 National Renewable Energy Laboratory
2 Verrus LLC

Produced under direction of Verrus LLC by the National Renewable Energy Laboratory (NREL) under TSA2431261.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

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

Strategic Partnership Project Report
NREL/TP-2C00-94844
June 2025

Vulcan Test Platform: Demonstrating the Data Center as a Flexible Grid Asset

Deepthi Vaidhynathan,[1] Kumaraguru Prabakar,[1] Gregory Martin,[1] Anand Ramesh,[2] Ben Wheeler,[2] Christopher Coco,[2] Jimmy Clidaras,[2] Matthieu Monsch,[2] Sangsun Kim,[2] Saurav Talukdar,[2] and Shilpa Marti[2]

1 National Renewable Energy Laboratory
2 Verrus LLC

Suggested Citation
Deepthi Vaidhynathan, Kumaraguru Prabakar, Gregory Martin, Anand Ramesh, Ben Wheeler, Christopher Coco, Jimmy Clidaras, Matthieu Monsch, Sangsun Kim, Saurav Talukdar, and Shilpa Marti. 2025. *Vulcan Test Platform: Demonstrating the Data Center as a Flexible Grid Asset*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-2C00-94844. https://www.nrel.gov/docs/fy25osti/94844.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

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

Strategic Partnership Project Report
NREL/TP-2C00-94844
June 2025

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 for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. Support for the work was also provided by Verrus LLC under TSA2431261. The views expressed in the article 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.

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

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover photos (clockwise from left): Josh Bauer, NREL 61725; Visualization from the NREL Insight Center; Getty-181828180; Agata Bogucka, NREL 91683; Dennis Schroeder, NREL 51331; Werner Slocum, NREL 67842.

NREL prints on paper that contains recycled content.

Disclaimer

This document and the information contained, referred to, referenced, or linked to herein is intended solely for informational and discussion purposes and is provided on an "as is" and "as available" basis. To the full extent permissible by law, Verrus disclaims all representations and warranties of any kind, express or implied, including, without limitation, any implied warranties of merchantability and fitness for a particular purpose. To the full extent permissible by law, Verrus will not be liable for any damages of any kind, including, any loss of profits, loss of use, business interruption, or indirect, special, incidental, consequential or punitive damages of any kind in connection with content, services, products or any other information provided or otherwise described herein. Verrus disclaims any obligation to provide updates to any information contained in this document or in connection with any related discussions.

Trademarks Verrus™, the stylized Verrus logo, Grid Aware, Carbon Aware, and Compute Aware™ are trademarks of Verrus LLC in the United States and/or other countries.

Copyright Verrus owns or licenses all copyright rights in all content, including, without limitation, all text, images, videos, and graphics in this document, to the full extent provided under the copyright laws of the United States and other countries. This copyright prohibits any act of copying, reproducing, modifying, distributing, displaying, performing or transmitting any of the content in this document for any purpose without the express consent of Verrus.

© 2025 Verrus LLC. All rights reserved.

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

Acknowledgments

The authors thank the National Renewable Energy Laboratory (NREL) team who supported the project and the technical report development. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. We used NREL's Advanced Research on Integrated Energy Systems (ARIES) virtual emulation environment. Specifically, we thank the following individuals and express our gratitude for your support:

Reviewers:

- Yaswanth Nag Velaga
- Vikram Roy Chowdhury
- Steven Hammond
- Benjamin Kroposki
- Katie Wensuc (editorial).

NREL supporters:

- Ty Ferretti
- Jennifer Kurtz
- Kristin Munch
- Andrew Hudgins
- Heather Newell
- Tyler Pinover
- Peter Gotseff
- Kurtis Buck
- Brian Fedish
- John Nangle
- Shwaathi Deenadayalan
- Owen Tanuwidjaja
- Steven Hammond
- Kristin Munch
- Barry Mather
- Marc Day
- Ray Grout
- Benjamin Kroposki
- Bill Tumas
- Juan Torres.

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

List of Acronyms

ARIES	Advanced Research on Integrated Energy Systems
BESS	battery energy storage system
CHIL	controller-hardware-in-the-loop
DNP3	Distributed Network Protocol 3
DRTS	digital real-time simulation
GFL	grid following
GFM	grid forming
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IT	information technology
NREL	National Renewable Energy Laboratory
PFMS	power flow management system
PUE	power usage effectiveness
SOC	state of charge
TCP	Transmission Control Protocol
UPS	uninterruptible power supply

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

Executive Summary

Increasing data center demand is outpacing grid infrastructure development. Artificial intelligence workloads and hyperscale cloud growth are creating unprecedented demand for power, while traditional grid expansion faces multiyear development timelines.

Verrus was founded by Sidewalk Infrastructure Partners to develop an innovative solution for this challenge: data centers that function as grid assets rather than passive loads. The Verrus approach integrates a novel grid-aware power flow management system with long-duration battery energy storage systems into a microgrid-controlled, medium-voltage power distribution architecture that delivers three critical capabilities:

- Fast-responding demand flexibility that can service requests from the utility within 10 s
- Uninterrupted transition to islanded operation during grid disturbances
- Continuous uptime assurance while maintaining all customer service level agreements.

Through Verrus's strategic partnership with the National Renewable Energy Laboratory (NREL), these capabilities were validated using NREL's Advanced Research on Integrated Energy Systems (ARIES) virtual emulation environment to model a 70-MW grid-interactive data center. We have successfully demonstrated that Verrus' data centers can deliver meaningful grid services without compromising mission-critical availability.

This technical report outlines the design, methodology, and results of this emulated deployment, demonstrating that data centers can play a pivotal role in enhancing grid flexibility and reliability without sacrificing service level guarantees.

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

Table of Contents

Disclaimer . . . iv
Acknowledgments . . . v
List of Acronyms . . . vi
Executive Summary . . . vii
1 Introduction . . . 2
2 Vulcan Test Platform Overview . . . 3
3 Data Center Power System . . . 5
4 Power Flow Management System . . . 9
5 Results . . . 13
5.1 Scenario 1: Utility Demand Response . . . 13
5.2 Scenario 2: Utility Outage, Islanded Operation, and Grid Reconnection . . . 14
6 Conclusion . . . 22
7 References . . . 23

List of Figures

Figure 1. Vulcan test platform: a CHIL-based framework for the validation of Verrus's data center power system operational sequences . . . 4
Figure 2. Conceptual single-line diagram of the data center power system . . . 8
Figure 3. Scenario 1: Utility demand response simulation in Vulcan . . . 14
Figure 4. Scenario 1: Utility demand response request for 35 MW of BESS discharge received (around $t = 304$ s) and the response of the PFMS dispatch of the BESS units in less than 5 s. . . . 15
Figure 5. Scenario 2, Step 1 and 2: Normal operation in grid-connected mode till $t < 180$ second. Grid voltage drops from 230 kV to 2 kV, the PFMS executes then executes the islanding algorithm and transitions the power system from grid connected mode to islanded mode using BESS units provide the backup power by operating in GFM mode. . . . 17
Figure 6. Scenario 2, Step 3, 4 and 5: BESS units continue to provide back up power with decrease in SOC. At $t \sim 940$ second, when average SOC of BESS units reach below 50%, the request to shed low priority/ interruptible workloads is sent to IT workload owner. Their action is simulated using a multiplier $M = 0.7$ . . . 18
Figure 7. Scenario 2, Step 5: Increase in estimated runtime on BESS due to 30 %reduction in load. 19
Figure 8. Scenario 2, Step 6: The BESS units reach the low SOC threshold of 15%, the PFMS bring generators online to support the load. . . . 19
Figure 9. Scenario 2, Step 7: Grid voltage is restored to 230 kV. The PFMS reconnects the power system back to the grid. The interruptible workloads return around $t = 1620$. . . . 20
Figure 10. Scenario 2, Step 8: The BESS units starts to charge after reconnection to the grid. . . . 21

List of Tables

Table 1. Scenario 1 Sequence of Events . . . 13
Table 2. Scenario 2 Sequence of Events . . . 15

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

1 Introduction

The U.S. power grid faces increasing challenges as data center power demand increases across the country. With artificial intelligence workloads, cloud expansion, and digital transformation driving multi gigawatt interconnection queues, utilities and grid operators are experiencing difficulties in meeting demand. Traditional solutions, such as peak load support plants and transmission upgrades, are often taking longer than expected to complete.

Among these constraints, a promising alternative is emerging: **flexible demand**. A 2025 report by Duke University's Nicholas Institute [1] estimates that modest curtailment of only 0.25% to 1% of annual consumption during critical hours could enable the grid to absorb up to 100 GW of new load without the need for new generation or other major upgrades.

Verrus was founded on the belief that data centers can deliver this kind of flexibility—**not as a compromise to availability, but as a new standard for resilience**. The Verrus data center design represents a fundamentally different architectural approach by integrating a grid-aware control plane, a battery energy storage system (BESS), and intelligent power flow optimization across data center campuses. Rather than relying on diesel generators to provide grid relief, this design leverages BESS for firm dispatch to the grid that can be relied upon for resource adequacy planning. Not only do BESS enable meaningful load reductions during times of utility duress, but also they can do so with faster response times than conventional diesel-based systems, providing an additional element of grid stability.

This report introduces **Vulcan**, a 70-MW grid-interactive data center test platform developed in collaboration with the National Renewable Energy Laboratory (NREL) and powered by NREL's Advanced Research on Integrated Energy Systems (ARIES) virtual emulation environment. Vulcan incorporates a digital real-time simulation of Verrus data center design and integrates it with Verrus' intelligent power flow management management system using a microgrid controller-hardware-in-the-loop (CHIL) framework. This enables rigorous testing of grid-interactive control strategies, real-time dispatch coordination, and system resilience behaviors under realistic electrical conditions. Vulcan was built to address the questions being raised by system operators, planners, and policymakers today: *How much can flexible loads actually help? Under what conditions? What capabilities must data centers develop to be part of the solution?*

In the following sections, we explore these questions in depth, and we offer a blueprint for the next generation of grid-interactive infrastructure that is ready for operational deployment today.

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

2 Vulcan Test Platform Overview

Verrus' data center design replaces traditional diesel generators with a BESS as the primary backup power source for both facility and information technology (IT) workloads during utility outages. Integrated with a 23-kV medium-voltage distribution system and advanced power flow optimization, the BESS also enables grid-interactive capabilities, such as demand response—transforming the data center from a passive load into an active, flexible asset that supports grid stability during peak demand periods. The medium-voltage power distribution architecture is agnostic to the type of backup power, and, as such, additional generators, fuel cells, or other units can be deployed as secondary backup to ensure that critical loads are maintained during extended grid outages. The medium voltage is converted to low voltage (415/480 V AC) through a step-down transformer to supply power to IT and associated cooling loads. Appropriate redundancies exist in the medium-voltage and low-voltage aspects of the power system to allow system availability higher than 99.99+ %. The system allows concurrent maintainability in all aspects of the electrical and mechanical infrastructure, with no single points of failure.

Given the high uptime requirements, along with the goal of providing services to the grid, a dedicated control system is needed to manage the power flow. Verrus has developed a novel control system, referred to as the power flow management system (PFMS), that delivers operational data center functionality while enabling the data center to respond to requests from the grid. The PFMS design is inspired by the Institute of Electrical and Electronics Engineers (IEEE) open standards, such as IEEE 2030.7 [2] and IEEE 1547-2018 [3]; and it interfaces with/controls power hardware, such as BESS, generators, and switchgear relays. The PFMS also interfaces with higher-level systems, including supervisory control and data acquisition and tertiary control systems, such as Verrus' optimal dispatch engine, but these are not a focus of this report and hence are shown as dotted interfaces in Figure 1.

The complexity, scale, and novelty of the PFMS and Verrus' data center power system calls for a rigorous validation and de-risking of the integrated power system before field deployment. With this intent, Verrus collaborated with NREL to develop the Vulcan test platform, which uses a CHIL-based validation of the PFMS integrated with a digital real-time simulation (DRTS) of the data center power system, as shown in Figure 1. A CHIL framework enables the validation of the sequence of operations for a wide collection of system states and operating scenarios. The DRTS can model and simulate power system components in the electromagnetic transient domain in real time. This allows DRTS users to perform a suite of experiments that is typically called a CHIL framework [4]. A CHIL framework leverages the ability of the DRTS to model electromagnetic transient models and run them in real time to generate power system states, such as voltage and current signals (analog), circuit breaker status/control (digital), and the PFMS response/control actions in the form of set points/ discrete inputs. The CHIL setup can be used to leverage the analog and digital input-output through network-based open communication protocol cards on the DRTS platform to execute closed-loop experiments with the PFMS [5].

To evaluate the PFMS in a CHIL framework, the data center power system is modeled in a DRTS platform, and the PFMS is connected to the DRTS model using Distributed Network Protocol 3 (DNP3). This setup constitutes a closed-loop framework and allows for the evaluation of the sequence of operations in the PFMS device. In summary, the Vulcan test platform comprises the following:

1. The data center power system, which is being modeled using the DRTS at NREL (part of NREL's ARIES virtual emulation environment). The data center power system is described in detail in Section 3.

2. A real-time automation controller using DNP3 connected as a communication device between the PFMS and the DRTS

3. The Verrus-developed PFMS, which is integrated with the real-time automation controller and uses DNP3 as the communication protocol to exchange data. Further details on the PFMS and its algorithms are described in Section 4.

This report documents how the Vulcan test platform demonstrates a grid-required demand response event using the BESS (referred to as Scenario 1) and power system stable operation during the loss of utility

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

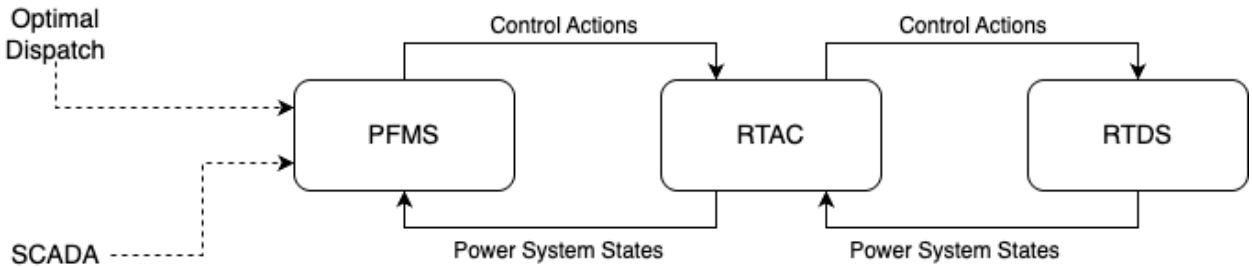

Figure 1. Vulcan test platform: a CHIL-based framework for the validation of Verrus's data center power system operational sequences

power (referred to as Scenario 2). Through these scenarios, we demonstrate that BESS units can provide support to the grid or respond to grid outages. Compute workloads are *not* impacted, thus validating the case for developing grid-interactive data centers. Detailed descriptions of these scenarios and test results are discussed in Section 5.

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

3 Data Center Power System

This section describes the data center power system that is modeled in a DRTS platform, along with the data objects that represent critical components for the control of the power flow. A simplified representation of Verrus's power system is shown in Figure 2, which includes two utility feeds (with breakers $X_{u,1}$ and $X_{u,2}$) connected to a medium-voltage distribution system that provides the power to the downstream low-voltage IT/computing servers and associated cooling loads. The medium-voltage power distribution has appropriate redundancies in the distribution network design to meet the needs of a high-uptime data center. For this report, only one utility feed is used in normal operation. The feed associated with $X_{u,1}$ is active, and the feed associated with $X_{u,2}$ is unused. The state of the grid, $Grid.State$, is considered stable if the voltage and frequency are within acceptable thresholds, and unstable otherwise. Additionally, if the utility initiates a demand response event, then $Grid.DR$ is considered $true$, and $false$ otherwise.

1. $Grid.State \in \{Stable, Unstable\}$
2. $Grid.DR \in \{true, false\}$

The aggregate capacity allocated to the IT loads is 50 MW, and the capacity allocated for the cooling loads is 20 MW on a shared 70-MW power domain. The individual power domains for the IT and compute loads connect to the medium-voltage power distribution through normally closed breakers $X_{IT,1}$,$X_{IT,2}$,...,$X_{IT,p}$. Note that a typical data center design would have a redundant power feed for the IT loads, but we have not incorporated that because faults/failures are not a focus of this report. The cooling infrastructure/ power domains have mechanical loads such as pumps and chillers, and are connected to the medium-voltage power distribution through the breakers $X_{c,1}$,$X_{c,2}$,...,$X_{c,q}$. These breakers are normally closed to provide power to the cooling systems. The IT load domains are modeled using a library load component with active power values from an open-source dataset from Google [6]. Each IT load domain, i, can be described as:

$$IT_i.P = M * P_i(k), i = \{1, 2, \ldots, p\}$$

where $IT_i.P$ is the active power of the i-th IT load domain, $P_i(k)$ represents the time-stamped values obtained from the Google dataset, and M is a multiplier percentage (taking the value of 0% to 100%) to model the optional load shedding initiated by the data center tenant (workload owner). $M = 100\%$ means NO load shedding, and $M = 0\%$ means full load shedding. Any intermediate value of M is to be interpreted as shedding the compute load of lower priority. Hyperscale data centers already have priority-based load shedding/reduction implemented [7], [8], and in practice the manifestation of the load shedding would be a signal to the IT workload owner, and they would be the one to curtail the lower-priority IT workloads. In this report, we use the multiplier M to simulate the actions of the IT workload owner to execute the load shedding. The total IT load, $IT.P$, is given as:

$$IT.P = \sum_i IT_i.P, i \in \{1, \ldots, p\}$$

The IT load domains have in-line uninterruptible power supply (UPS) to provide momentary power during grid outage events until backup power is restored. The UPS is modeled as a voltage source with a switching logic to transition the load between stable power on the feeder or the UPS. The total cooling load, $Cooling.P$, is modeled assuming a constant power usage effectiveness (PUE), that is:

$$Cooling.P = (PUE - 1) * IT.P$$

This study uses a PUE of 1.3. The total cooling load is assumed to be split equally across the q cooling domains. The $j-th$ cooling domain's active power is given by:

$$Cooling_j.P = Cooling.P/q, j = \{1, 2, \ldots, q\}.$$

Note that constant PUE modeling is assumed for simplicity of the study. In a real scenario, PUE is a function of environmental variables, such as temperature and humidity, which have temporal variations. Additionally, cooling power allocation is not equally spread across the cooling load domains, but for the

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

simplicity of the study, it is assumed that the load is equally spread. Power consumption of the cooling air and water movement in the server halls is not considered because it typically represents a small fraction ($< 3\%$) of the total power consumption. The cooling loads are NOT supported by UPS, and will go through a restart during utility outage and transition to back up power. The total system load, P_{Load}, in the power system is $IT.P + Cooling.P$. The IT and cooling load domains are modeled as a $P-Q$ loads in the DRTS.

The BESS units, as shown in Figure 2, are rated at 4.5 MW (P_{BESS}^{limit}) with an energy capacity of 18 MWh. These units are capable of operating in grid-following mode (GFL) and grid-forming mode (GFM). The BESS units connect to the medium-voltage distribution system through breakers $X_{b,1}$,$X_{b,2}$,...,$X_{b,m}$. The BESS breakers are nominally closed. The aggregate BESS inverter capacity is 81 MW. For a duration of 4 hours, the aggregate energy capacity is 324 MWh for the entire data center building; thus, the BESS is sized to support the entire 70 MW of rated IT and cooling loads. Each BESS unit, $i \in \{1,2,\ldots,m\}$, is represented as a data object, $BESS_i$:

- $BESS_i.SOC \in [0,100]\%$
- $BESS_i.P$
- $BESS_i.V$
- $BESS_i.I$
- $BESS_i.f$
- $BESS_i.P^* \in [0,4.5]MW$
- $BESS_i.V^* \in [90,110]\% \times V_{nominal}$
- $BESS_i.f^* \in [98,102]\% \times f_{nominal}$
- $BESS_i.Mode \in \{GFL, GFM, Idle\}$,
- $BESS_i.Available \in \{True, False\}$

where $BESS_i.SOC, BESS_i.P, BESS_i.V, BESS_i.I, BESS_i.f$ denotes the state of charge (SOC), active power, voltage, current, and frequency of the i-th BESS, respectively. The operating range for $BESS_i.V^*$ and $BESS_i.f^*$ is based on common practice, which can differ across manufacturers. $BESS_i.P^*$ is the active power command/set point for the $i-th$ BESS when the operating mode $BESS_i.Mode = GFL$. When the operating mode of the $i-th$ BESS is GFM, the voltage set point and frequency set point are denoted as $BESS_i.V^*$ and $BESS_i.f^*$, respectively, and they are specified as a percentage of the nominal voltage ($V_{nominal} := 480$ V) and frequency ($f_{nominal} := 60$ Hz). The idle mode represents a state where the BESS unit is ready to operate in either GFM or GFL mode and is awaiting a command from the PFMS. The BESS unit could be available if it has sufficient stored energy, or it could be unavailable for use if the SOC is too low or due to any other internal fault condition of the unit. The availability aspect of the unit is defined by $BESS_i.Available$ field. The BESS units are modeled using a DRTS library model.

Although any type or count of backup power units can be added to the Verrus medium-voltage power domain (up to the rating threshold of the medium-voltage power infrastructure), for demonstration purposes, the representation in the Vulcan model is a full complement of standby diesel generators (70 MVA). The diesel generator units, rated at 3.5 MVA each (P_{gen}^{limit}), connect to the medium-voltage power distribution through breakers $X_{g,1}$,$X_{g,2}$,...,$X_{g,n}$. The aggregate generator capacity is 77 MW, and it has the redundant generator unit capacity baked in. The diesel generator units are used only during grid outages after the energy stored in the BESS has depleted to a minimum threshold; thus, the generator capacity is sized to support the entire 70 MW of the data center IT and cooling loads. Each diesel generator, $i \in \{1,2,\ldots,n\}$, is represented as:

$$Gen_i.P^* \in [30,100]\% \times P_{gen}^{limit},$$
$$Gen_i.Available \in \{True, False\}$$

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

where $Gen_i.P^*$ denotes the active power set point of the i-th generator, and $Gen_i.Available$ denotes whether that unit is available to serve the load. The generator units are modeled using a DRTS library model.

All BESS, IT, cooling, and generator breakers described here are controllable with power monitoring capability, and they are represented as:

$$
\begin{aligned}
&X_{a,b}.State \in \{Open, Close\} \\
&X_{a,b}.P \\
&X_{a,b}.V \\
&X_{a,b}.f \\
&X_{a,b}.I
\end{aligned}
$$

where $X_{a,b}$ denotes any generic breaker mentioned here; and V, P, f, I denote the voltage, active power, frequency, and current measured at the breaker. The $X_{a,b}.State$ denotes the open or closed status, which can be controlled by the PFMS.

This study focuses on the control of the active power in the power system; hence, for a simpler presentation, we do not define the reactive power representations for these loads and sources. In this power system, the BESS units are to be operated in GFL mode when the data center is grid connected, and they would be used to provide a reduction in the demand of utility power from the data center by discharging the BESS upon request of the utility. For the rare cases of a loss of utility power, the BESS units would be operated in GFM mode to provide stable backup power to the data center loads. For longer-duration utility outages, the generators would ramp up as the stored energy of the BESS units deplete. When the utility power returns, the data center would reconnect to the power grid, and the BESS would recharge to its target SOC. The PFMS, described in the next section, is responsible for coordinating the various assets in the data center power system to accomplish the desired operational behavior described here.

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

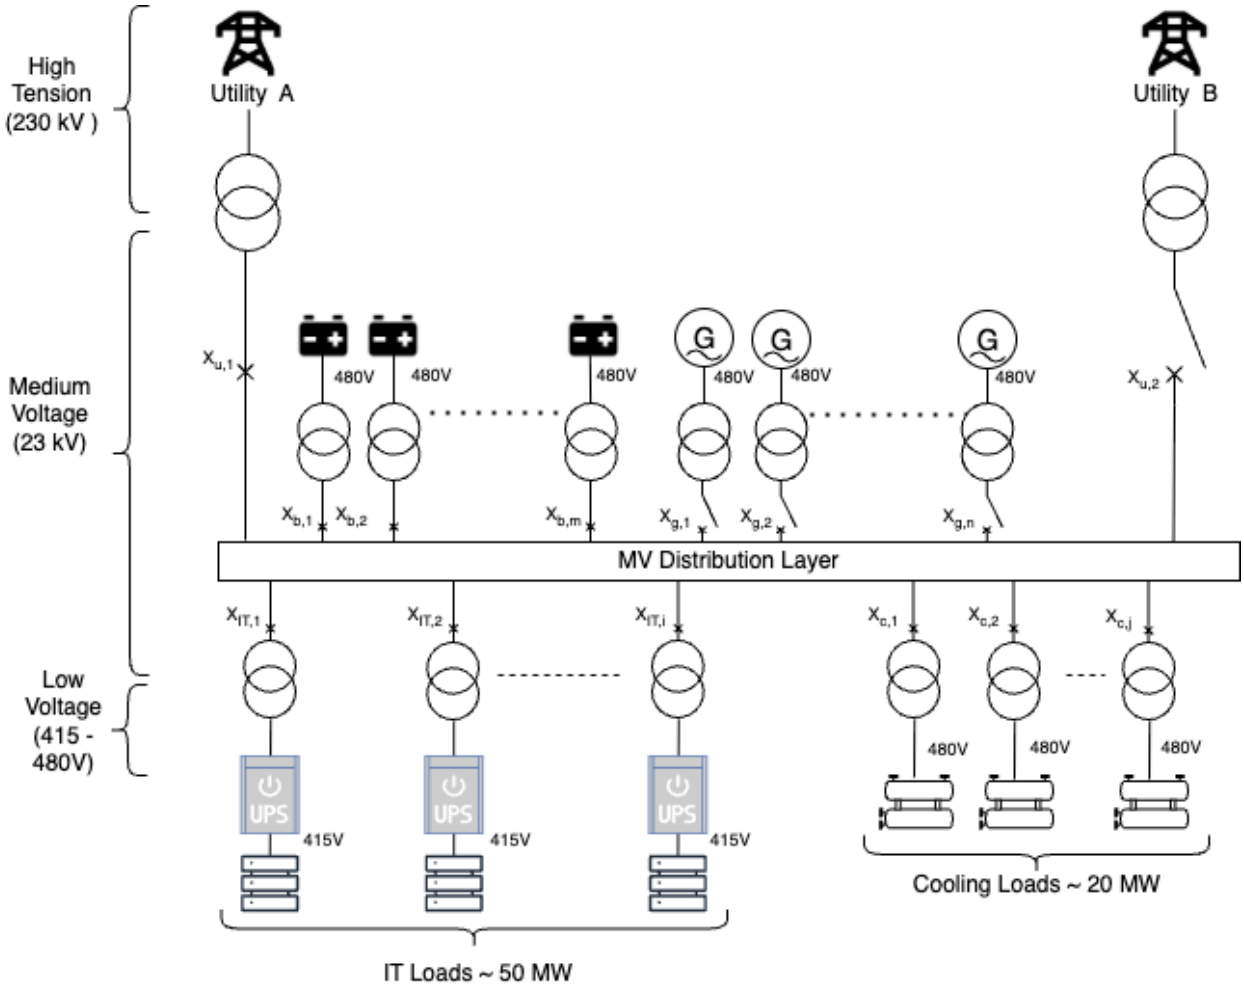

Figure 2. Conceptual single-line diagram of the data center power system

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

4 Power Flow Management System

The PFMS is a secondary control system responsible for ensuring stable, uninterrupted power delivery to the loads in the data center and for using the data center assets, such as the BESS, to respond to requests from the grid. The PFMS is critical to ensuring the stability of the data center power system while improving the resilience of the grid by responding to utility dispatch requests. The stable power delivery to loads requires the PFMS to constantly monitor the status of the grid, BESS, generators, and loads, and decide on the control actions/sequences of operation to be executed. The design of the PFMS is based on open standards, such as IEEE 2030.7 and IEEE 1547-2018; and it uses standard communication interfaces, such as Modbus Transmission Control Protocol (TCP), DNP3, and International Electrotechnical Commission (IEC) 61850. It monitors and controls power hardware, such as the BESS, generators, and relays; and it can also interface with higher-level systems, such as supervisory control and data acquisition systems and advanced distribution management systems. For this study, we focus on the PFMS control actions for the following:

- Responding to a grid request of active power dispatch from the BESS units
- Isolating the data center power system from the grid (also referred to as islanding) in case of unstable grid conditions
- Reconnecting the data center to the grid when the voltage or frequency disturbance has subsided.

For simplicity, note that in this study we do not consider reactive power dispatch, but it can be included in the same framework. In addition, generators are not used to respond to the dispatch requests from the utility, but rather they are used by the PFMS solely to provide backup power to the data center in case of an extended outage (after the BESS units have reached their predefined minimum SOC). Next, we define algorithms for the PFMS to accomplish these scenarios by testing on the Vulcan platform. Note that these algorithms are not the only way to accomplish the desired functionalities, and Vulcan provides the flexibility to test other algorithms in the future.

1. Grid-connected operation: When the grid voltage and frequency are within an acceptable range (as defined by IEEE 1547-2018), the PFMS operates the data center with power from the grid and keeps $X_{u,1}$ closed. The PFMS keeps the BESS units charged to a target *SOC* (SOC_U) to have sufficient capacity to provide backup power to the data center in case of a grid outage or to respond to requests from the utility for demand response. The same active power set point (P_{charge}) (charging is considered negative, and discharging is considered positive) is applied to all the available BESS units, and it is capped by the physical upper limits of charging (P^{limit}_{charge}). P_{charge} is computed based on the available headroom from the power system physical limit ($Facility^{limit}$) to support the total load (P_{Load}) with some buffer (P^{buffer}). The description is summarized in Algorithm 1.

Algorithm 1 Grid-Connected Operation

$X_{u,1} \leftarrow Close$
$P_{charge} = Facility^{limit} - P_{Load} - P^{buffer}$
for $i = 1, ..., m$ **do**,
  if ($BESS_i.Available == True$) and ($BESS_i.SOC < SOC_U$) and($Grid.DR == False$) **then**
    $BESS_i.Mode \leftarrow GFL$
    $BESS_i.P^* \leftarrow -\min(P_{charge}/m, P^{limit}_{charge})$
  end if
end for

2. Grid-connected operation with utility-requested BESS dispatch: When the data center is operating in grid-connected mode, the PFMS keeps the BESS units charged (with P_{charge} as the charging power) to the target upper *SOC* limit (SOC_U), and it discharges the available BESS units (N) to meet the aggregate dispatch (ρ_d) requested by the utility. When the utility requests a dispatch from the PFMS, *Grid.DR* is set to *true*, and *false* otherwise. The PFMS commands equal active power

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

dispatch from the available BESS units to balance the run time across the units. Upon discharging, when the BESS units reach the low SOC limit (SOC_L), they are deemed unavailable and placed in idle mode by the PFMS. Once the utility-requested dispatch event concludes, the PFMS resumes normal grid-connected operation and charges the BESS units to the target SOC. The PFMS control action in the form of a pseudo code is described in Algorithm 2.

Algorithm 2 Grid-Connected Operation with Utility-Requested BESS Dispatch

if $Grid.DR == True$ **then**
  for $i = 1, ..., m$ **do**,
    if $BESS_i.Available == True$ **then**
      if (($BESS_i.SOC > SOC_L$) and) **then**
        $BESS_i.Mode = GFL$
        $BESS_i.P^* \leftarrow \min(P^{limit}_{discharge}, \rho_d/N)$
      else if $BESS_i.SOC < SOC_L$ **then**
        $BESS_i.Mode \leftarrow Idle$
      else
        $BESS_i.P^* \leftarrow 0$
      end if
    end if
  end for
else
Continue grid-connected operation
end if

3. Grid-connected to off-grid/islanding process: The PFMS monitors the voltage, frequency, and rate of change of frequency of the incoming power from the grid and compares that against the minimum and maximum allowable thresholds with the associated time durations. If the grid voltage, frequency, or rate of change of frequency are outside the allowable range for a sustained duration, the PFMS confirms the criteria for grid power being unstable as being met and disconnects the data center by opening breakers $X_{u,1}$ and $X_{u,2}$. For this study, we implemented the thresholds and time durations as defined in IEEE 1547-2018. After disconnecting the data center power system from the utility, the PFMS then commands the available BESS units to initiate a black-start process in GFM mode to provide stable nominal voltage and frequency in the islanded power system. In this phase, the IT load is powered by the UPS while the cooling loads (not backed by the UPS) shut down. Once the BESS restores the nominal voltage and frequency on the medium-voltage power distribution with a sufficient number of units online to support the data center load, the UPS units transfer the IT load to the BESS, and the cooling systems restart. The islanding process is summarized in Algorithm 3.

4. Off-grid/islanded operation: In the case of a longer-duration outage from the grid, the data center power system will deplete the energy stored in the BESS units. As the median BESS SOC drops below a defined threshold SOC_{LS} ($SOC_U > SOC_{LS} > SOC_L$), the PFMS would create alerts for the data center IT workload owner that the remaining energy on the BESS is limited, and the operator could consider shedding less critical compute workloads to increase the run time on the BESS. To simulate the action of the IT operator shedding lower-priority compute workloads, in this study, the multiplier of $M = 0.7$ is used when the median SOC across the BESS units reaches below SOC_{LS}. At this stage, only critical workloads are being supported, and shedding noncritical workloads would extend the run time on the BESS. As the SOC of the BESS units drops below SOC_L, the PFMS would consider that BESS unit as unavailable for use and bring an appropriate number of generator units online to meet the critical load demand while accounting for the load increases in the form of reserve capacity, $P_{reserve}$. The off-grid/islanded operation is summarized in Algorithm 4.

5. Off-grid to grid-connected/reconnection process: When the utility power has returned and the voltage and frequency are within the acceptable range for a sufficient time duration, the PFMS would

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

Algorithm 3 Islanding Process

if islanding criteria $==$ *true* **then**,
  $Grid.State \leftarrow Unstable$
  $X_{u,1}.State \leftarrow Open, X_{u,2}.State \leftarrow Open$,
  for $i = 1, 2, \ldots, m$ **do**
    if ($BESS_i.SOC > SOC_L$) and ($BESS_i.Available == true$) **then**,
      $X_{b,i}.State \leftarrow Close$
      $BESS_i.V^* \leftarrow V_{nominal}$
      $BESS_i.f* \leftarrow f_{nominal}$
      $BESS_i.Mode \leftarrow$ GFM Mode
    else
      $BESS_i.Available \leftarrow False$
    end if
  end for
else
Continue grid-connected operation
end if

Algorithm 4 Islanded Operation

BESS units in GFM support P_{Load}
if $median([BESS_i.SOC]) < SOC_{LS}$ **then**,
Send alert to IT operator to consider shedding load to increase remaining run time of the BESS units,
  $M \leftarrow 0.7$,
else
  for $i = 1, 2, \ldots, m$ **do**
    if ($BESS_i.SOC < SOC_L$) and ($BESS_i.Available == true$) **then**,
      if $\sum_{k,k\neq i} BESS_k.P < P_{reserve} + P_{Load}$ **then**
        $Generator_{count} = ceiling[(P_{reserve} + P_{Load} - \sum_{k,k\neq i} BESS_k.P)/P_{gen}^{limit}]$
        $X_{g,r}.State \leftarrow Close, r = \{1, \ldots\ldots, Generator_{count}\}$,
        $Gen_r.P^* \leftarrow P_{gen}^{limit}, r = \{1, \ldots, Generator_{count}\}$,
        $BESS_i.Mode \leftarrow Idle$
      end if
    end if
  end for
end if

consider the grid reconnection criteria (sometimes referred to as the enter service criteria) as met and would initiate the reconnection process. In this study, we adopted the thresholds and time duration as specified in IEEE 1547. We adopt an open transition approach [9] to transition the power system from operating with backup power sources to operating with power from the grid. Once the reconnection criteria are met, the PFMS first disconnects the loads from the medium-voltage distribution system and then turns OFF/disconnects the generators and commands the BESS into idle mode. It then closes the breakers $X_{u,1}$ and, finally, closes all the load breakers. Once the grid power is restored in the power system, the PFMS commands the BESS units to charge if their *SOC* is below SOC_U by changing the mode to *GFL*. The reconnection is described in Algorithm 5.

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

Algorithm 5 Reconnection Process

if reconnection criteria $== true$ **then**,
$Grid.State \leftarrow Stable$
$X_{IT,i}.State \leftarrow Open, X_{c,j}.State \leftarrow Open$, for all $i \in \{1,\ldots,p\}, j \in \{1,\ldots,q\}$,
$Gen_k.Mode \leftarrow OFF$, for all $k \in \{1,\ldots,n\}$
$X_{g,k} \leftarrow Open$, for all $k \in \{1,\ldots,n\}$
$BESS_l.Mode \leftarrow Idle$, for all $l \in \{1,\ldots,m\}$
$X_{u,1} \leftarrow Close$
$X_{c,j}.State \leftarrow Close$, for all $j \in \{1,\ldots,q\}$,
$X_{IT,i}.State \leftarrow Close$, for all $i \in \{1,\ldots,p\}$,
$BESS_l.Mode \leftarrow \{GFL\}$
Continue grid-connected operation
else
Continue islanded operation
end if

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

5 Results

In this section, we review the results from the integration of a real-time simulation model of Verrus's data center power system and controller hardware for the PFMS in the Vulcan test platform. We show that Verrus's data center power system can provide grid support when requested by the utility while maintaining the high uptime requirements of the IT loads. We illustrate this claim through the following two scenarios.

5.1 Scenario 1: Utility Demand Response

Table 1 shows the steps of the demand response scenario that are simulated in the Vulcan test platform. The process starts with the grid providing 100% of power to the IT and cooling loads and the BESS being idle at the target SOC of 80 % (Step 1). Next, 35 *MW* of BESS dispatch is requested by the utility (Step 2). The PFMS then provides the requested BESS dispatch (Step 3, Algorithm 2). Once the demand response event ends, the PFMS starts to charge the BESS (Step 4) to the target *SOC* of 80%(Step 5). We show four plots in Figure 3 for Scenario 1: (a) the utility demand response request, (b) the grid and BESS power, (c) the average SOC of the BESS units, and (d) the IT and cooling load.

Table 1. Scenario 1 Sequence of Events

Step No.	State	Description
1	Grid-connected operation	Loads are 100% powered by the grid.
2	Demand response event starts	Utility issues demand response call for 35 MW of dispatch for 35 min (shortened from 4–8 h for the simulation)
3	BESS dispatch	BESS discharge at 35 MW, loads powered by utility and BESS
4	Demand response event ends	Return to normal grid-connected operation, BESS starts charging
5	Grid-connected operation	100% grid powered and BESS recharged to target *SOC*

Figure 3(a) shows the utility demand response request, where 1 indicates that the data center is being requested by the utility for an active power dispatch from the BESS units, and 0 means there is no BESS dispatch request from the utility. The utility demand response is *true* from (time denoted as t) $t = 5$ min to $t = 35$ min, during which the utility-requested active power dispatch from the BESS units is $35MW$; outside this window, there is no utility dispatch request.

When $t < 5$ min: Figure 3(a) shows that the demand response request is false; and Figure 3(b) and Figure 3(d) show that the entire load is supported by the grid, and the BESS power output is zero. The *SOC* in Figure 3(c) is constant at the target *SOC* of 80%(using Algorithm 1).

When $5 < t < 35$ min: In Figure 3(a), the utility demand response request for 35 *MW* of BESS active power dispatch is received around $t = 5$ min, and the PFMS responds to this request (using Algorithm 2). The BESS total output increases to 35 *MW* within 5 s (see Figure 4 for a zoomed-in view around the $t = 5$-min or 300-s time window), and the power drawn from the grid decreases accordingly, as shown in Figure 3(b). In this time frame, the average *SOC* of the BESS units is decreasing linearly in Figure 3(c) due to the constant power output from the BESS units. Although the power drawn from the grid decreases substantially, the total IT and cooling loads do not experience any interruption, as no effect is observed on the load profiles in Figure 3(d). The PFMS helped reduce the effective data center load on the grid to cater to the utility's request during the peak demand hours and thus improved the grid resilience.

When $t > 35$ min: The utility demand response ends, as shown in Figure 3(a). The PFMS starts to charge the BESS to the target the *SOC* (Algorithm 2) with 15 MW of charging power, which explains the negative BESS output power curve shown in Figure 3(b). The charging of the BESS can also be verified by the increasing SOC in Figure 3(c) for $t > 35$ min. The total IT and cooling load, along with the additional load due to charging the BESS, is supported by power from the grid. Accordingly, the power drawn from the grid increases compared to the grid power drawn for $t < 5$ min. Throughout Scenario 1, the IT and compute loads do not experience any power-related interruptions while the energy mix changes between grid power and the BESS units to meet the load requirements during the demand response.

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

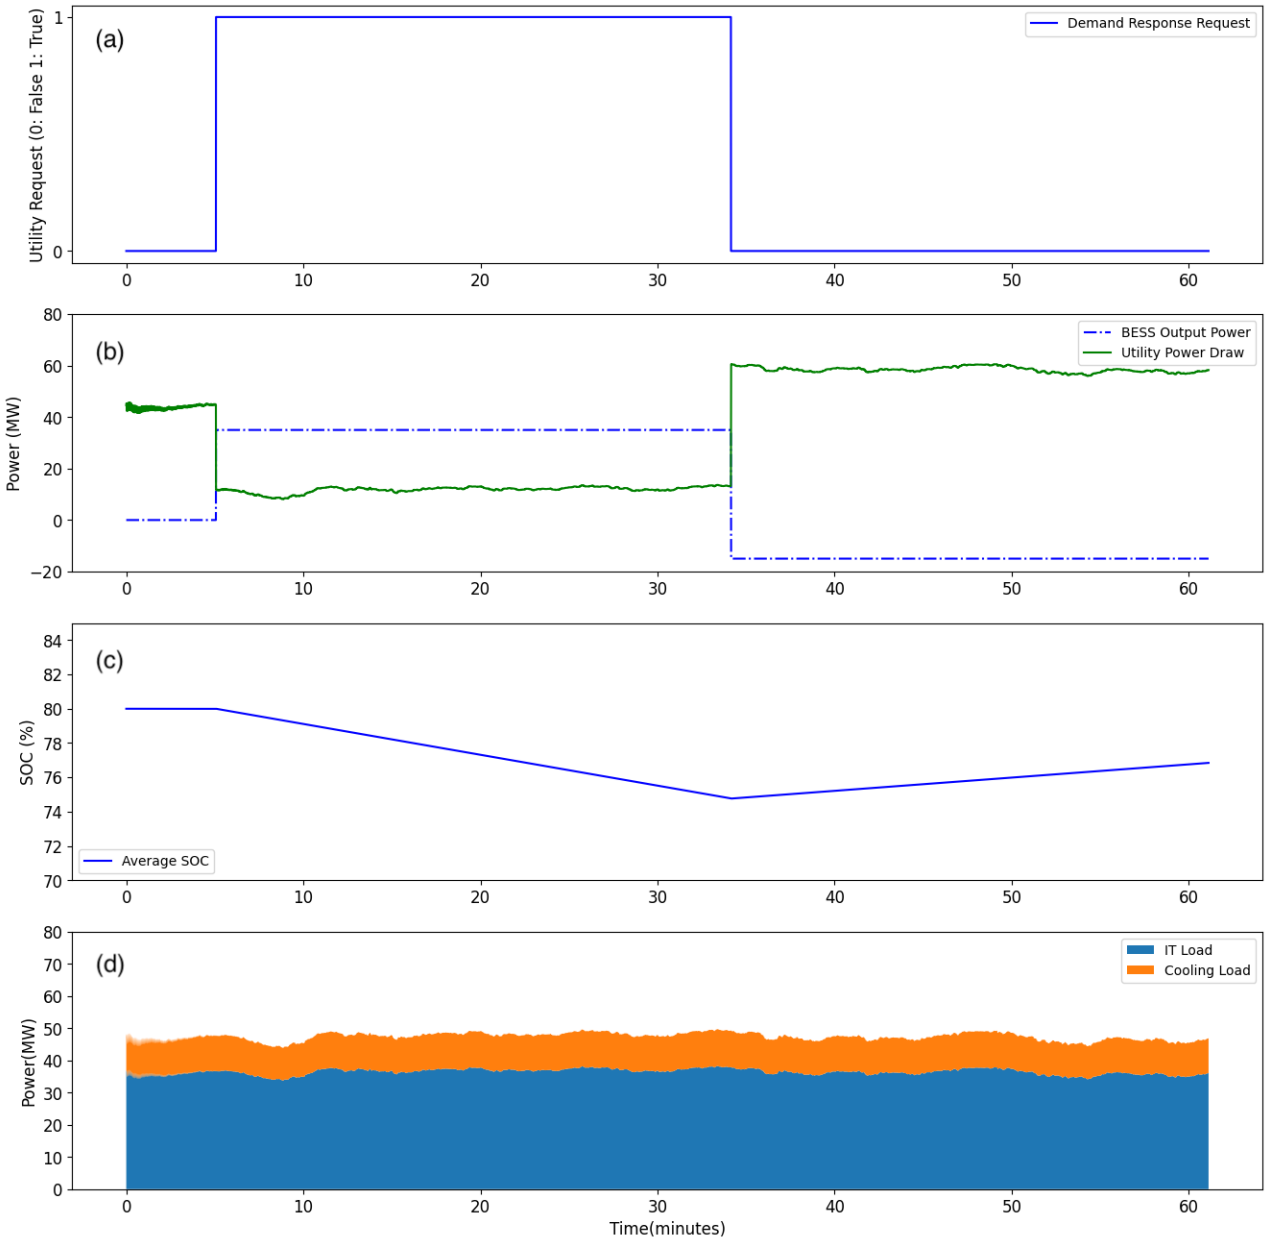

Figure 3. Scenario 1: Utility demand response simulation in Vulcan

5.2 Scenario 2: Utility Outage, Islanded Operation, and Grid Reconnection

This scenario is focused on the power system islanding after a grid power outage and reconnection to the grid when a stable grid is available. First, the PFMS uses the BESS units to provide backup power, and then, as the energy stored in the BESS units depletes and they reach below the low *SOC* threshold, it makes use of the generators. Once the power provided by the grid is stable, the PFMS will reconnect the system to the grid and initiate charging of the BESS units.

Table 2 describes the specific steps involved in this scenario that are simulated in the Vulcan test platform. The first step is the initial state of the system, where the grid is stable and provides power to all the loads in the data center. In Step 2, there is an outage in grid power (simulated by a grid voltage drop from 230 kV to 2 kV), which triggers the islanding process in the PFMS (using Algorithm 3). In this open transition, the IT loads are momentarily supported by UPS, following which all loads are supported

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

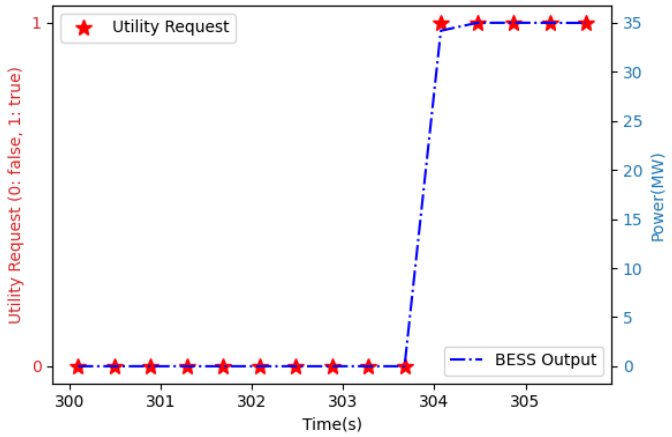

Figure 4. Scenario 1: Utility demand response request for 35 MW of BESS discharge received (around $t = 304\ s$) and the response of the PFMS dispatch of the BESS units in less than 5 s.

Table 2. Scenario 2 Sequence of Events

Step No.	State	Description
1	Grid-connected operation	100% grid powered
2	Utility outage	Transition to islanded mode and 100% powered by BESS after UPS transition
3	Full IT workload running on BESS	Depletion of BESS *SOC*
4	(Optional) load-shedding request to IT operator	Request via secure, open standard communication protocol (OpenADR [10])
5	Low-priority IT workload shed	Extends BESS run time
6	Dispatch long-duration power backup (generators)	Continue uninterrupted uptime for mission-critical IT workloads
7	Utility recovery	Return to grid-connected operation, interruptible workloads return
8	Grid-connected operation	BESS starts to charge

entirely by the BESS units operating in GFM. In Step 3, there is continued power system operation in islanded mode (using Algorithm 4), which results in a decrease in the *SOC* of the BESS units. At this stage, the IT workload owner could consider shedding low-priority/interruptible workloads to preserve the energy stored in the BESS units and extend their run time. This is described in Step 4. The PFMS is capable of interfacing with the IT workload owners to send the load-shedding requests (through open communication interfaces, such as OpenADR). Step 5 simulates the load shedding of the interruptible workloads with a multiplier, M, of 0.7. Eventually, the *SOC* of the BESS units drops below the low threshold of 15 %, and the generators are brought online by the PFMS to meet the load demand (Algorithm 4). Finally, in Step 7, the grid is stable, and the PFMS reconnects the power system with the grid through an open transition process (using Algorithm 5). Once the power system is reconnected to the grid, the interruptible workloads return, which is simulated using $M = 1$. In the last step, the PFMS starts to charge the BESS to the target *SOC*, and the load is supported entirely by the grid (Algorithm 1).

Next, we discuss the results from the execution of Scenario 2 on the Vulcan test platform. We show four plots in each figure, plot (a) shows the grid voltage and plot (b) shows the power from energy sources (utility, BESS, Generators and UPS). In plot (c) the average SOC of the BESS units are shown and plot (d) shows the total IT and cooling load.

The initial state of the power system is grid-connected mode with 100% of the load supported with power from the grid, as shown in Figure 5 (a), when time $t < 180$ second. The BESS and the generator aggregate output is zero, and the *SOC* of the BESS is flat (because it is not discharging). Around 180 second mark, we simulate a voltage sag event and drop the grid voltage from 230 kV to 2 kV as shown in Figure

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

5 (a). The PFMS detects its islanding criteria to be met, and it executes the islanding process (see Algorithm 3). As shown in Figure 5, between $t = 180$ and $t = 210$, the power from the grid drops to zero after the utility breaker is opened by the PFMS. At this stage, there are no energy sources powering the power system, and the IT load is supported by the UPS. Next, PFMS uses the BESS units in GFM mode to provide backup power to the IT and cooling loads, as confirmed by the BESS output power in the Figure 5 (b) after $t = 210$ seconds. The average SOC of the BESS units start to decrease. Overall, it took less than 30 seconds to island the 70 MW power system after detecting the unstable grid conditions.

The power system continues to operate in off grid/ islanded mode with backup power provided by the BESS units. In Figure 6, as the SOC of the BESS drops below the threshold of 50 %(SOC_{LS}), the PFMS alerts the IT workload owner (via communication interfaces such as OpenADR) to consider shedding the low-priority/interruptible workloads to extend the run time of the BESS (Algorithm 4). We simulate a load shed around $t \sim 940$ second by using a multiplier of $M = 0.7$, as shown in Figure 6 (d). The load drops from approximately ~ 45 MW to ~ 35 MW, and the SOC curve of the BESS units shows the slowdown of the BESS depletion through the change of slope around the inflection point (just after the $t \sim 940$ second mark). Figure 7 shows the runtime remaining on the BESS increases by about 30 % due to shedding of interruptible workloads. Here, the runtime remaining is computed using a ratio of the remaining energy in the BESS units (based on average SOC of the units) to the last 5 second average load.

As the SOC of the BESS units drop further and reaches below the minimum SOC threshold of 15%(SOC_L), the PFMS brings online an appropriate number of generators to support the load (Algorithm 4). As shown in Figure 8 (b), the output of the BESS units drop around $t \sim 1120$ second, the generator units come online at the same time and deliver an aggregate output power of 44 MW. Since the generator output is higher than the loads, excess power of about 10 MW is used for charging the BESS units. This is seen by the increase in SOC in Figure 8 (c). As seen in Figure 8 (d), the loads do not experience any impact as the energy source is changing (between the BESS and generators).

Next, stable utility power returns, which is simulated by increasing the utility voltage back to 230 kV at $t \sim 1510$ second, as seen in Figure 9(a). The PFMS initiates the grid reconnection process to reconnect the power system to the grid using an open transition (Algorithm 5) after the reconnection criteria is met around $t \sim 1565$ second. As shown in Figure 9 (b), the loads are disconnected at $t \sim 1565$ second, and the compute/IT workloads are supported by the UPS. The PFMS then reduces the aggregate generator power output to zero around $t \sim 1565$ and the generators are turned off, as shown in Figure 9(b). Next, the PFMS puts the BESS units into idle mode, and their aggregate output is shown to drop to zero as well. At this point, there are no energy sources supplying power to the loads, and the compute/IT load are supported by the UPS. The PFMS then closes the utility breaker at $t \sim 1600$ second, and the power supplied by the grid is shown to increase as it supports the total load in Figure 9(b). The interruptible workloads return, which is simulated by the multiplier $M = 1$ around $t = 1620$ second, and is confirmed by the total load increasing from 35 MW to approximately 45 MW in Figure 9(d). Overall, it took less than 30 seconds to reconnect the power system to the power grid after confirming the stable grid conditions. Finally, the PFMS commands the BESS units to GFL mode and start charging at $t \sim 2180$ second, which is verified by the increase in SOC in Figure 10(c) and by the negative output power of the BESS in Figure 10(b) .

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

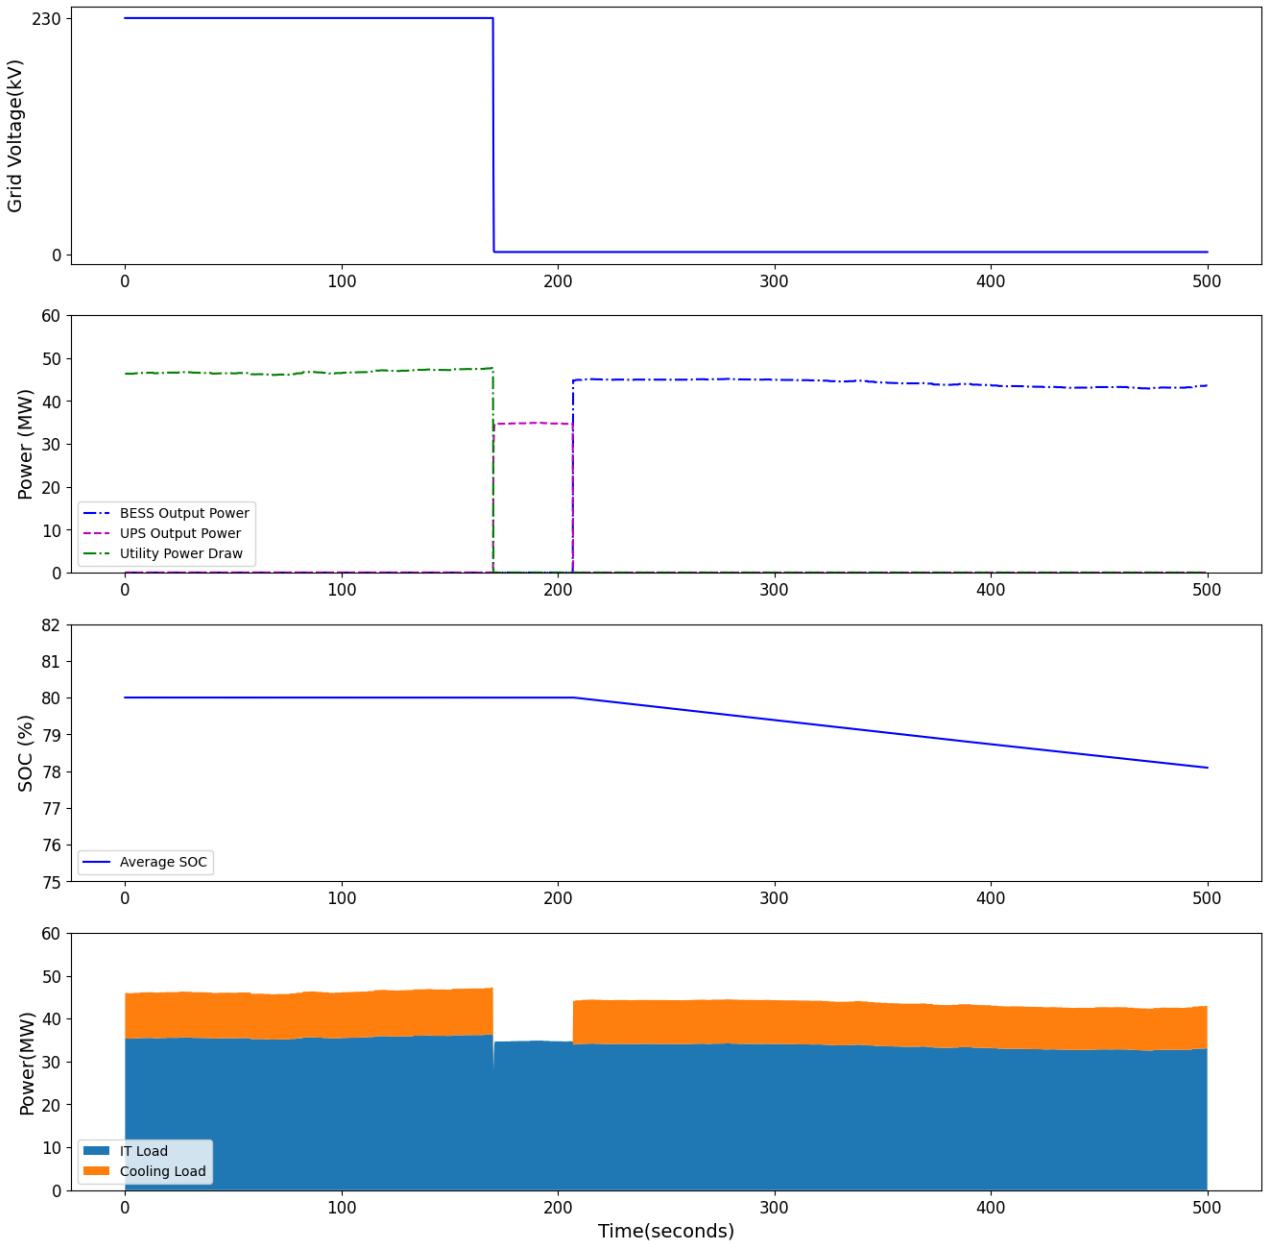

Figure 5. Scenario 2, Step 1 and 2: Normal operation in grid-connected mode till $t < 180$ second. Grid voltage drops from 230 kV to 2 kV, the PFMS executes then executes the islanding algorithm and transitions the power system from grid connected mode to islanded mode using BESS units provide the backup power by operating in GFM mode.

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

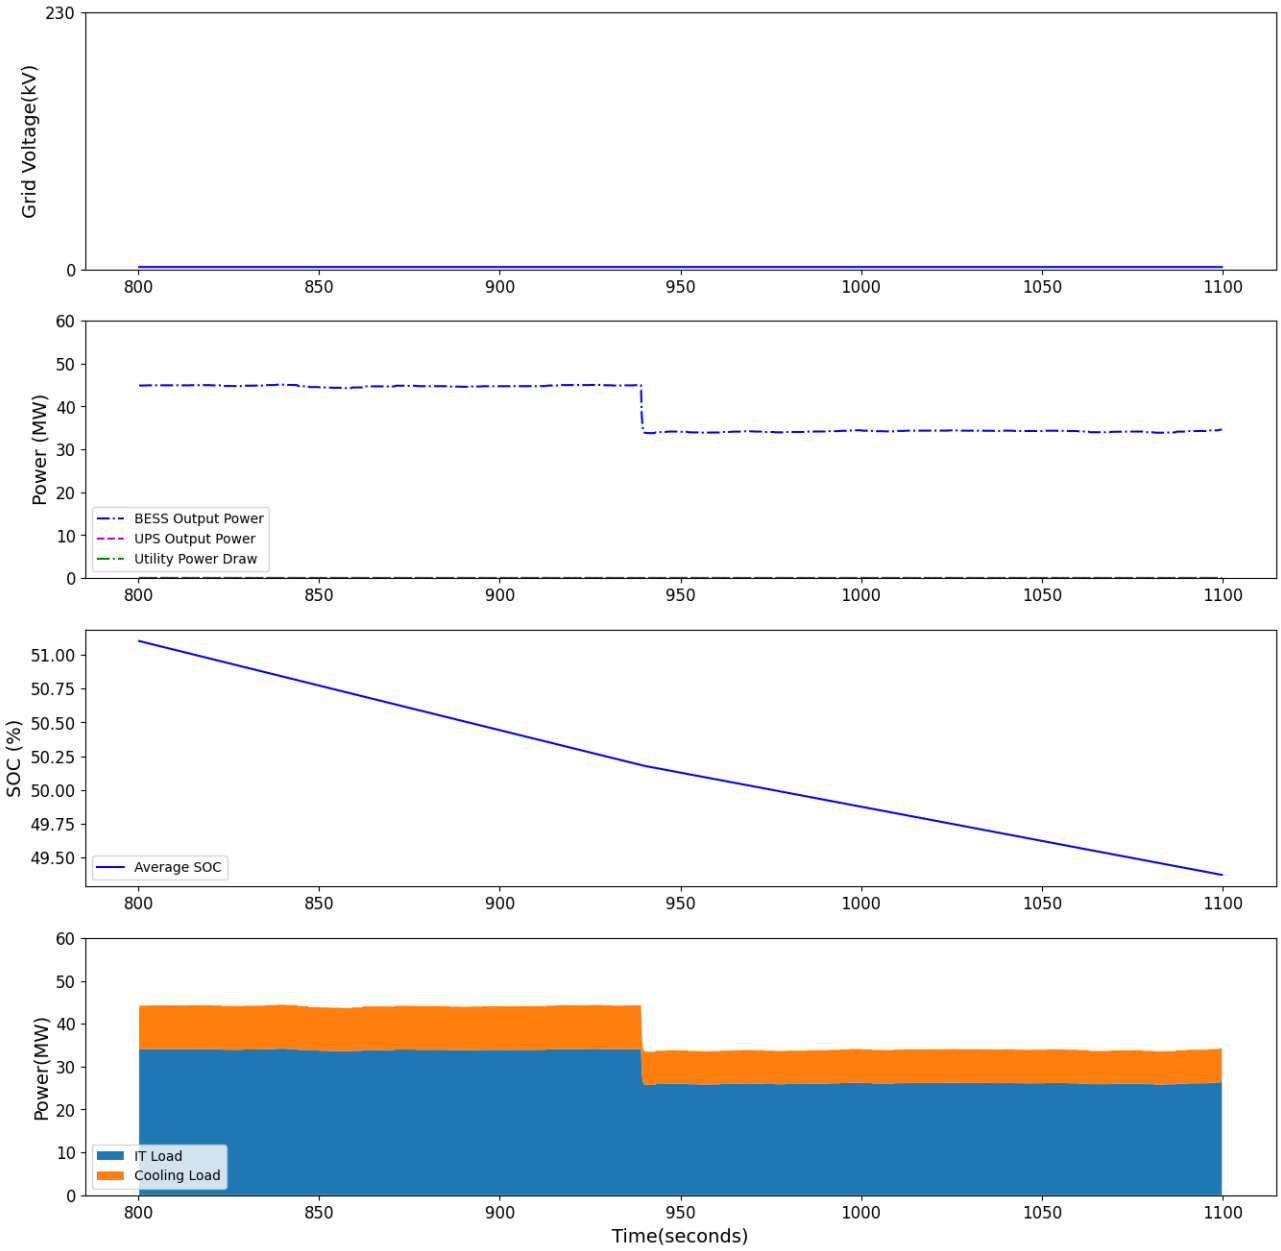

Figure 6. Scenario 2, Step 3, 4 and 5: BESS units continue to provide back up power with decrease in SOC. At $t \sim 940$ second, when average SOC of BESS units reach below 50%, the request to shed low priority/ interruptible workloads is sent to IT workload owner. Their action is simulated using a multiplier $M = 0.7$

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

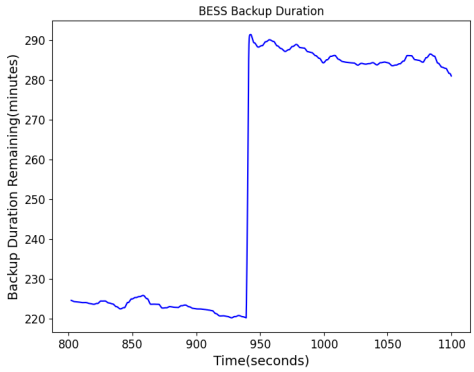

Figure 7. Scenario 2, Step 5: Increase in estimated runtime on BESS due to 30 %reduction in load.

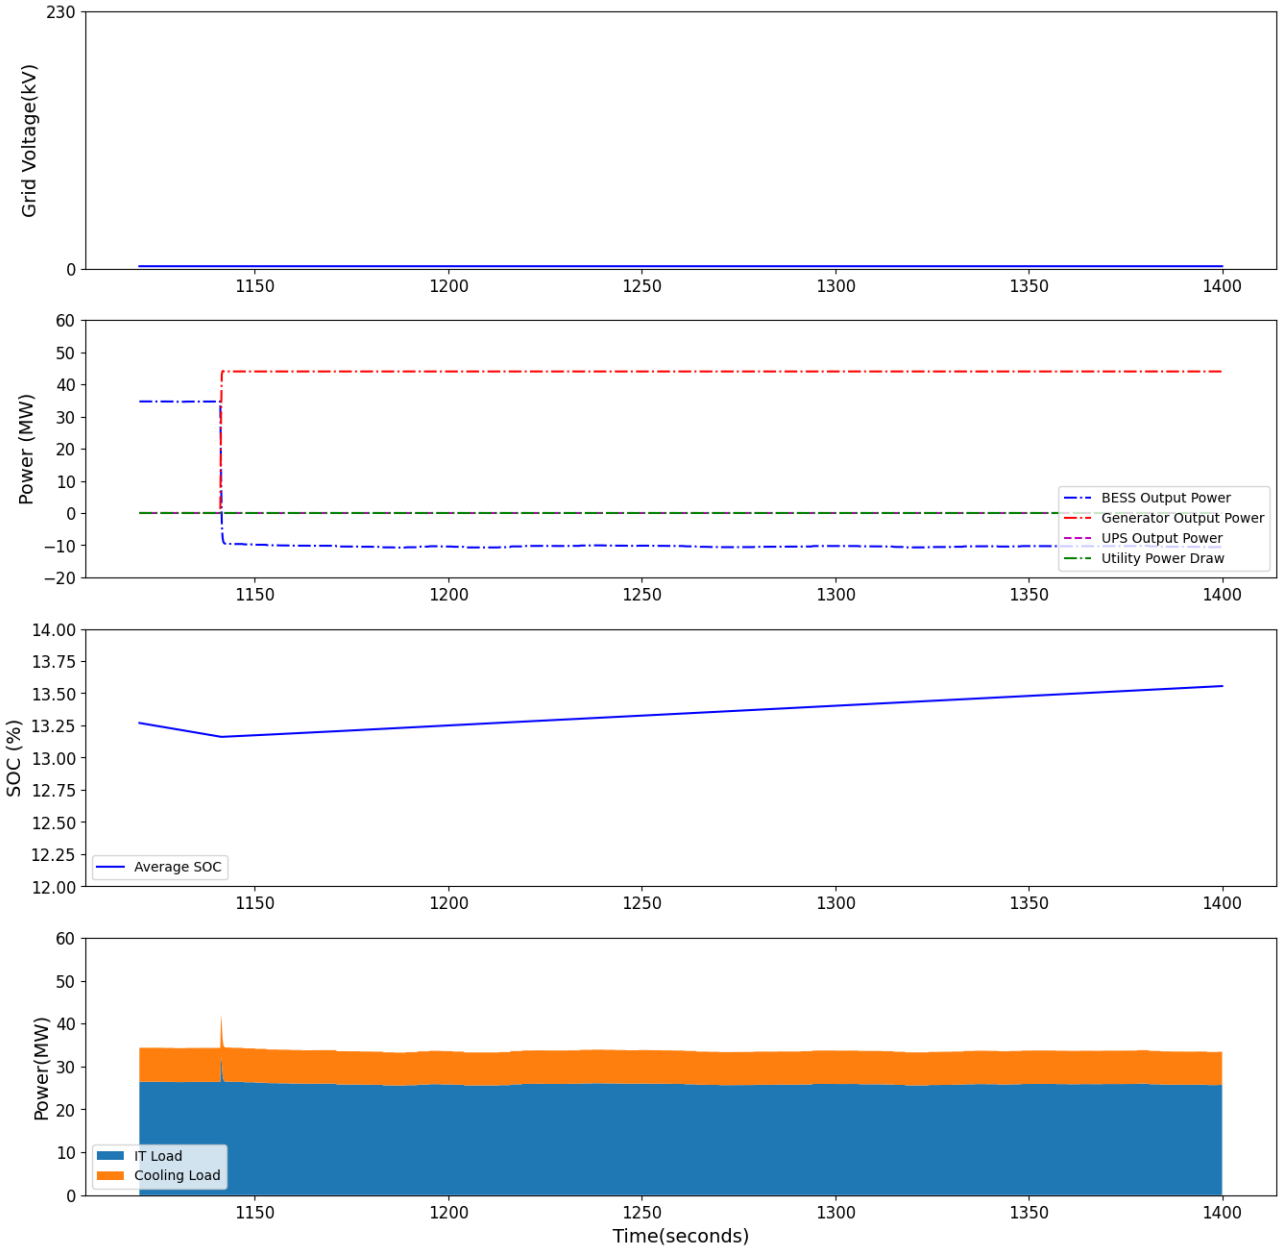

Figure 8. Scenario 2, Step 6: The BESS units reach the low SOC threshold of 15%, the PFMS bring generators online to support the load.

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

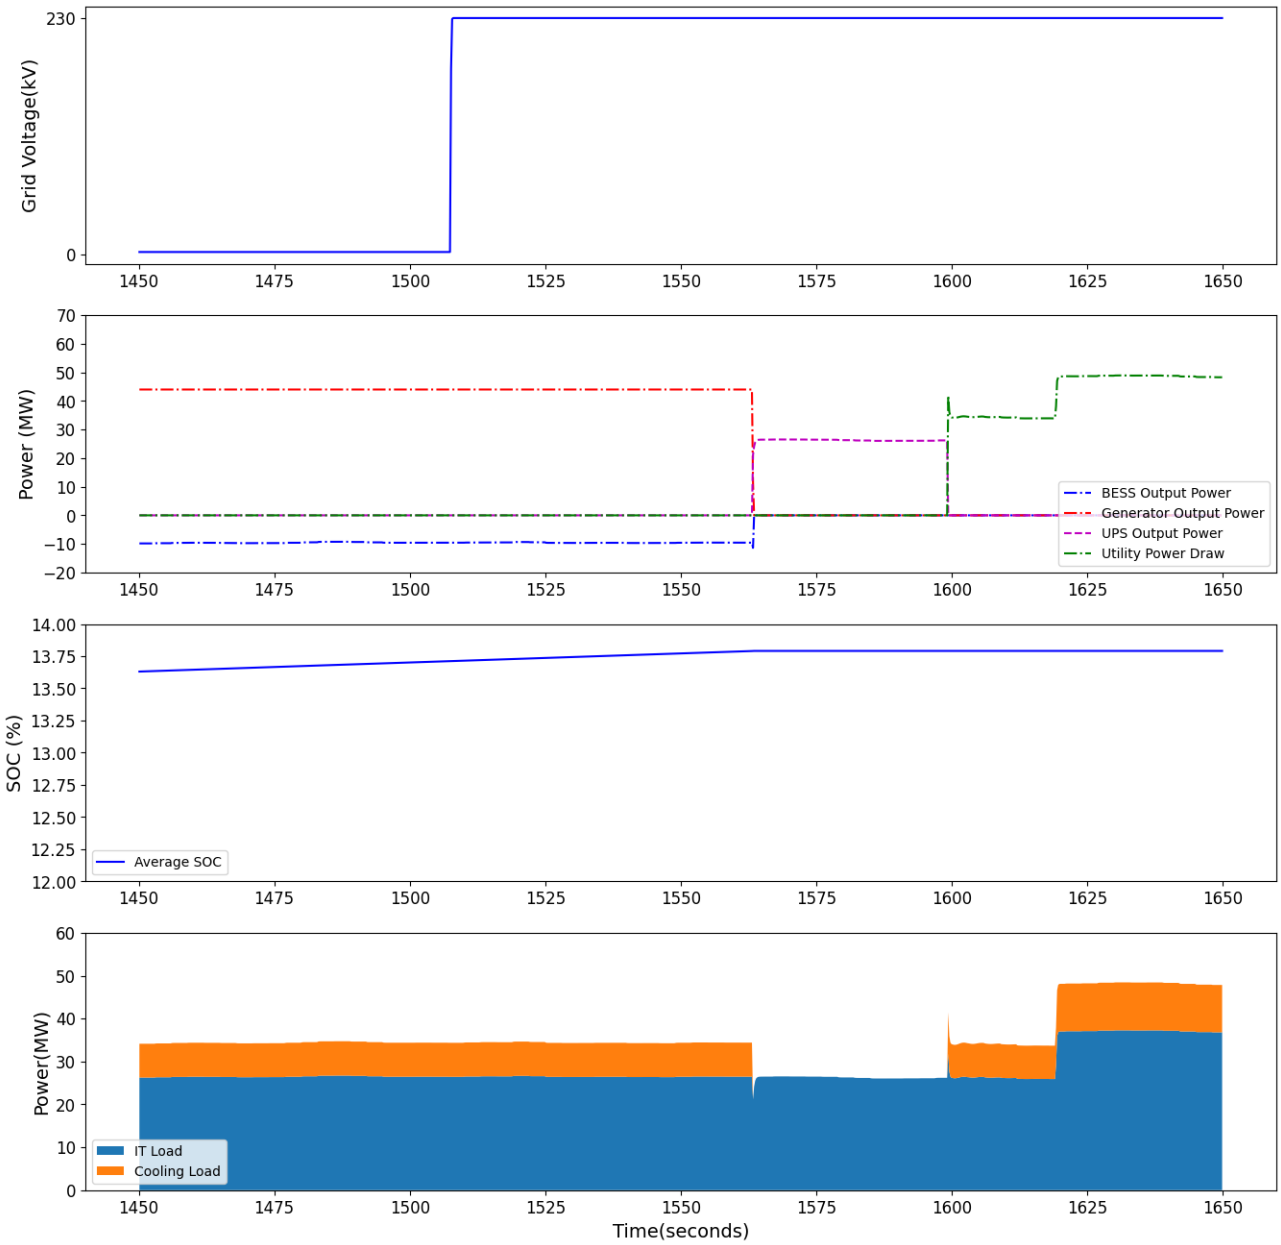

Figure 9. Scenario 2, Step 7: Grid voltage is restored to 230 kV. The PFMS reconnects the power system back to the grid. The interruptible workloads return around $t = 1620$.

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

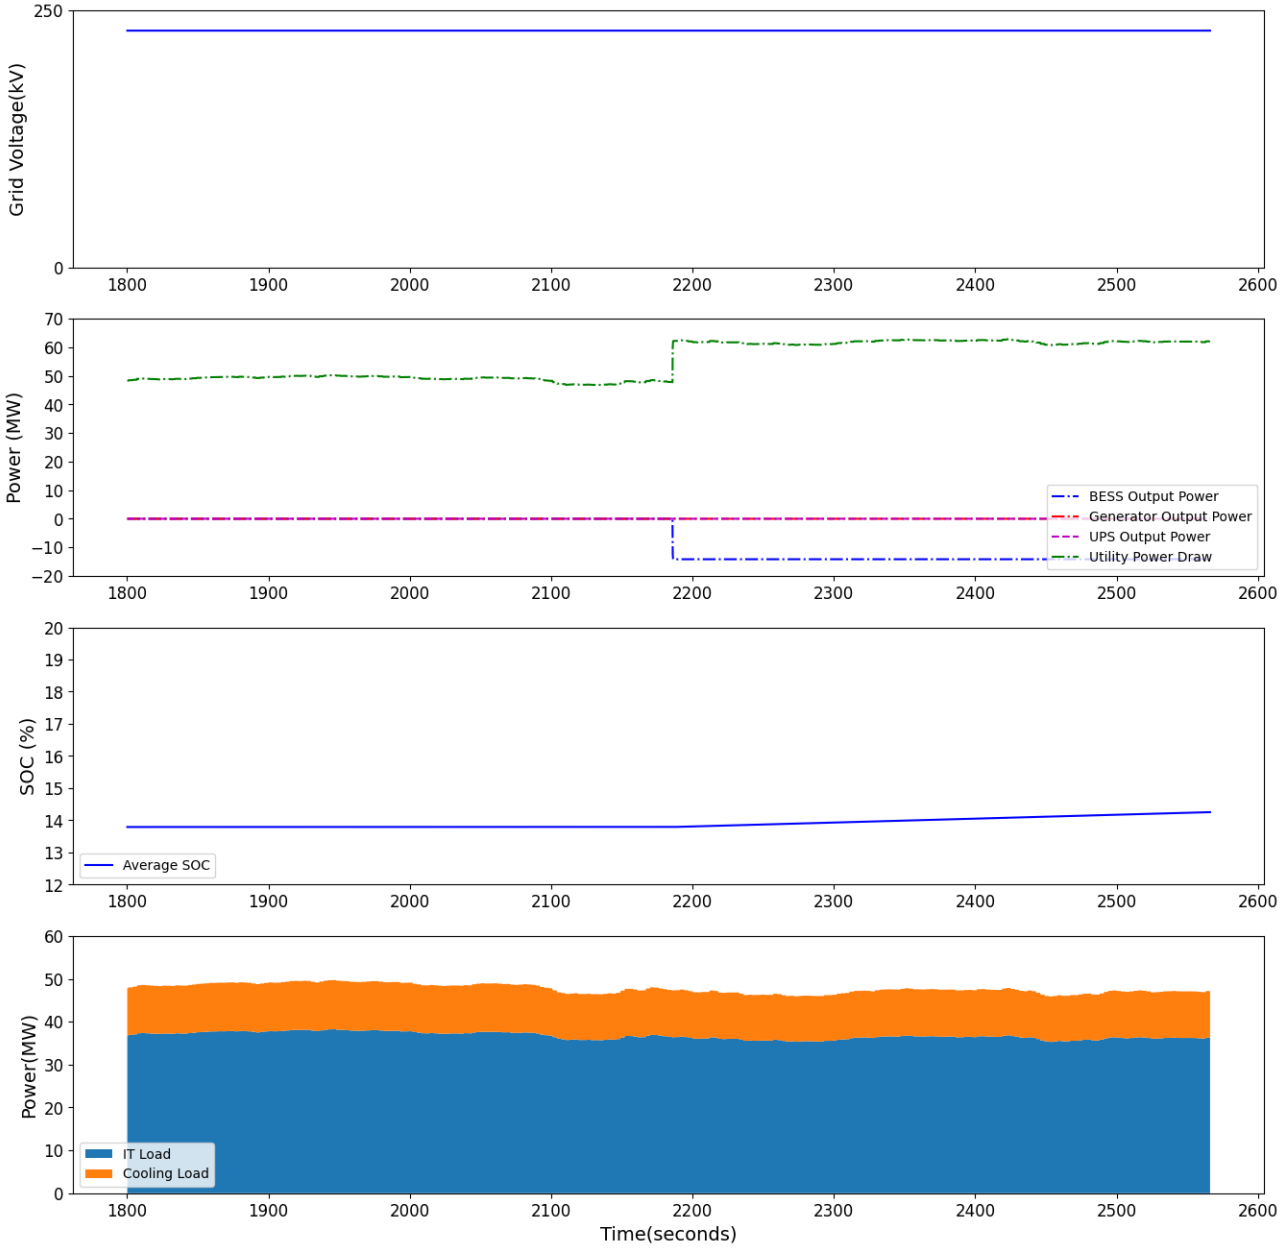

Figure 10. Scenario 2, Step 8: The BESS units starts to charge after reconnection to the grid.

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

6 Conclusion

Using the Vulcan test platform, we demonstrated that data centers can act as a grid asset by dispatching on-site BESS to support utility grid events. Additionally, we showed that during unstable utility conditions, the Verrus secondary control system (PFMS) can switch nearly instantaneously to the BESS units to provide backup power to the data center. The PFMS can also reconnect the data center to the grid once stable utility conditions return. The Vulcan test platform results confirm that the IT load is not impacted and continues to operate within the customer uptime commitments. Beyond the scenarios presented in this report, the Vulcan platform can be used for future case studies on grid disturbances and dynamic IT load profiles, and it is particularly well suited to study transient effects during interactions between the data center and the grid.

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

7 References

[1] T. Norris, T. Profeta, D. Patino-Echeverri, and A. Cowie-Haskell, *Rethinking load growth: Assessing the potential for integration of large flexible loads in us power systems*, 2025.

[2] "Ieee standard for the specification of microgrid controllers," *IEEE Std 2030.7-2017*, pp. 1–43, 2018. DOI: 10.1109/IEEESTD.2018.8340204.

[3] "Ieee standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," *IEEE Std 1547-2018 (Revision of IEEE Std 1547-2003)*, pp. 1–138, 2018. DOI: 10.1109/IEEESTD.2018.8332112.

[4] K. Prabakar, "Testing and evaluation methods for protection systems," in *Power System Protection in Future Smart Grids*, Elsevier, 2024, pp. 167–188.

[5] A. Maitra, A. Pratt, T. Hubert, D. Wang, K. Prabakar, R. Handa, M. Baggu, and M. McGranaghan, "Microgrid controllers: Expanding their role and evaluating their performance," *IEEE Power and Energy Magazine*, vol. 15, no. 4, pp. 41–49, 2017.

[6] J. Wilkes, *Yet more Google compute cluster trace data*, Google research blog, Posted at https://ai.googleblog.com/2020/04/yet-more-google-compute-cluster-trace.html., Mountain View, CA, USA, Apr. 2020.

[7] V. Sakalkar, V. Kontorinis, D. Landhuis, S. Li, D. De Ronde, T. Blooming, A. Ramesh, J. Kennedy, C. Malone, J. Clidaras, *et al.*, "Data center power oversubscription with a medium voltage power plane and priority-aware capping," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 497–511.

[8] Q. Wu, Q. Deng, L. Ganesh, C.-H. Hsu, Y. Jin, S. Kumar, B. Li, J. Meza, and Y. J. Song, "Dynamo: Facebook's data center-wide power management system," *ACM SIGARCH Computer Architecture News*, vol. 44, no. 3, pp. 469–480, 2016.

[9] S. Chakraborty, M. T. Rana, and M. V. Salapaka, "Active synchronization of islanded microgrid using droop-controlled grid-forming inverters," in *IECON 2022–48th Annual Conference of the IEEE Industrial Electronics Society*, IEEE, 2022, pp. 1–6.

[10] C. McParland, "Openadr open source toolkit: Developing open source software for the smart grid," in *2011 IEEE power and energy society general meeting*, IEEE, 2011, pp. 1–7.

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