5G Securely Energized and Resilient: (5G-SER): Final Report

Joshua Rivera, Brian Miller, Jordan Peterson, Erich Feth, Paul Snyder, and Tony Markel

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-5T00-88550
June 2024

5G Securely Energized and Resilient: (5G-SER): Final Report

Joshua Rivera, Brian Miller, Jordan Peterson, Erich Feth, Paul Snyder, and Tony Markel

National Renewable Energy Laboratory

Suggested Citation

Rivera, Josh, Brian Miller, Jordan Peterson, Erich Feth, Paul Snyder, and Tony Markel. 2024. *5G Securely Energized and Resilient: (5G-SER): Final Report*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5T00-88550. https://www.nrel.gov/docs/fy24osti/88550.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/TP-5T00-88550
June 2024

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

NOTICE

This work was authored 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 U.S. Department of Defense. 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.

Preface

This final report bridges the gap between the work completed in Task 3 of the 5G Securely Energized and Resilient (5G-SER) project and work completed with a Celona carrier-grade 5G network to run a series of grid scenarios using power hardware-in-the-loop.

For the 5G-SER project Tasks 2 and 3, NREL deployed an open-source 5G communications platform and built a distributed control system for grid-edge control via multi-access edge computing and an open 5G radio access network (Rivera et al. 2023). Task 3 grid infrastructure comprised simulated microgrid components running on a real time digital simulator. However, for the efforts of Task 4, we have replaced the real time digital simulator models with integrated power hardware-in-the-loop components to include a photovoltaic inverter, battery energy storage, and critical and noncritical load to execute the suite of tests from the previous tasks to revalidate physical device effectiveness using 5G wireless controls. In addition, Task 4 allowed us to upgrade system software components for the OpenAirInterface (OAI) 5G core, a cell tower henceforth referred to as a gNodeB, and user equipment integrated with the Celona 5G system. In Task 4, we also upgraded the distributed controls software to enable operational automation and grid resilience.

By incorporating the Celona 5G core network funded by another Office of the Under Secretary of Defense FutureG[1] project, we were able to cross-compare the OAI 5G core network performance with the commercial Celona 5G carrier-grade system performance.

[1] https://rt.cto.mil/futureg-home/

Acknowledgments

This work has been funded by the U.S. Department of Defense (DOD) Office of the Under Secretary of Defense for Research and Engineering as a complement to the FutureG initiative. We appreciate the forethought of Dr. Daniel Massey to engage in collaboration with the National Renewable Energy Laboratory and to initiate the project while recognizing the importance of energy and resilience in DOD's use of FutureG and the value of energy security and resilience. We also thank Dan Jasper for continuing to engage with the research team and linking us to relevant experts and opportunities for exposing the team's technical progress.

Many researchers and technical experts have contributed to this project during its 3-year duration. We acknowledge several that are not specifically part of the final project report authors: Daniel Bennett was instrumental in seeding the concepts for this work and executing our original collaboration agreement with DOD. Jared Macon successfully led and executed the work plans for the early years of the effort. Sukwha Kyung laid a foundation for the use of OpenAirInterface for an open-source radio access network, and Donald Price executed important tasks related to our frequency use plans and software-defined radio implementation.

List of Acronyms

5G	5th generation
5G-SER	5G Securely Energized and Resilient
5QI	5G traffic prioritization
DC	direct current
DCS	distributed control system
DER	distributed energy resource
GHz	gigahertz
gNB	gNodeB (or 5G Base Station)
HMI	human-machine interface
IED	intelligent electronic device
JSON	java script object notation
MEC	multi-access edge computing
MMS	manufacturing message specification
ms	millisecond
NREL	National Renewable Energy Laboratory
OAI	OpenAirInterface
PHIL	power hardware-in-the-loop
PV	photovoltaic
RAN	radio access network
UE	user equipment
UPF	user plane function
UPS	uninterruptable power supply
USRP	Universal Software Radio Peripheral

Executive Summary

The National Renewable Energy Laboratory (NREL) has completed work on 5G communications technology integration with physical power systems (versus simulated power systems demonstrated in prior work) under the 5G Securely Energized and Resilient program, thus successfully achieving a major technical milestone. Technical challenges were overcome to implement 5G end-to-end network environments with physical components, including microgrid switchgear, a grid-forming microgrid inverter, an energy storage battery, a solar PV inverter, a controllable load, and a critical load. In parallel, a distributed controls architecture for the microgrid components powering 5G network devices was also successfully modified from its previous instantiation for a simulated environment to work with these physical components. We then confirmed the feasibility of 5G wireless to enable resilient communications between controller and distributed solar and storage resources while exploring ways to configure 5G components to survive power disturbances. The results detailed in this report validated, through lab testing, prior simulated results that 5G wireless systems were able to provide resilient communication for the control of distributed microgrid power systems.

The work was conducted using both open-source and commercial 5G networking infrastructure. Early efforts implemented OpenAirInterface (OAI) systems, and mid-project we leveraged successful outcomes from other U.S. Department of Defense efforts to also integrate Celona 5G systems. As a result, we enabled a cross-comparison of the OAI system performance with a commercial carrier-grade Celona system using Federal Communications Commission licensure and spectrum allocation for full-power outdoor broadcasting in the Citizen's Broadband Radio Service frequency band n48. The results demonstrate that both OAI and Celona are viable options for performing research and development scenarios to understand the reliability and integration for controls and microgrid operations.

Prior work built the 5G networks and tested wireless control of simulated power components while the efforts addressed in this report highlight the outcomes of 5G integration and testing with physical power systems. The new test results validate our prior expectations and conclusions. As a result, the NREL team completed live outdoor field experiments and assessments for the Future Generation Wireless Technologies ("FutureG") program involving virtualized distributed edge controls for electric power systems. Using both open-source and commercial 5G cores (OAI and Celona, respectively), the team controlled a microgrid, which in turn powered the radio access network and core network devices as its critical load. Response time, power consumption, latency, and traffic prioritization under extreme network stress were observed successfully. Survivability of the communications network during a power outage was increased from hours to a duration of many weeks, and the recovery time for a controller crash or attack was reduced to minutes using automation and orchestration of communications components.

Table of Contents

List of Figures

List of Tables

1 Energy System Integration

The 5G Securely Energized and Resilient (5G-SER) final report is a culmination of work completed four phases. In Task 1, the team conducted preliminary research and designed a 5G microgrid test system to determine the feasibility of the program. For Task 2, the team deployed the emulated wireless microgrid and distributed controller components designed in Task 1 and developed a test plan to validate the resilience and robustness of the 5G microgrid. In Task 3, the team executed the test scenarios developed in Task 2 on an emulated microgrid. The results of the research for Tasks 1 through 3 were published in a publicly available report.[2]

Task 4 of the 5G-SER program expands on the energy system integration and controls progress achieved in prior simulated testing via the multi-access edge computing (MEC) and OpenAirInterface (OAI) 5G radio access network (RAN) with the real time digital simulator. Task 4 included the upgrade of OAI 5G systems and components, integration of the Celona 5G system, final design and prototyping of the distributed control system (DCS), and integration of all power hardware-in-the-loop (PHIL) components. Figure 1 represents the microgrid controls test bed used to perform all test cases and analysis. When preparing for this phase, the major finding was that 5G "ultralow latency" only applies to the final section of network, i.e., the MEC collocated at the gNodeB (gNB) cell tower, the radio link to the user equipment (UE), and the local distributed energy resource (DER) network connected to the UE. This finding was the prime reason why a distributed control approach was required, which led to our development of a custom DCS for testing.

Yellow	Represents administrative access to the MEC for hive node grid-edge control operations via the 5G system.
Blue	Represents the general packet radio service tunneling protocol user plane connection from the 5G core user plane function (UPF) through the gNB to UE on the workstation (this can also be described as the N3 interface)
Green	Represent the distributed controls communications for the worker node and the UE to the grid edge as well as through the 5G system out the UPF to the MEC hive node.

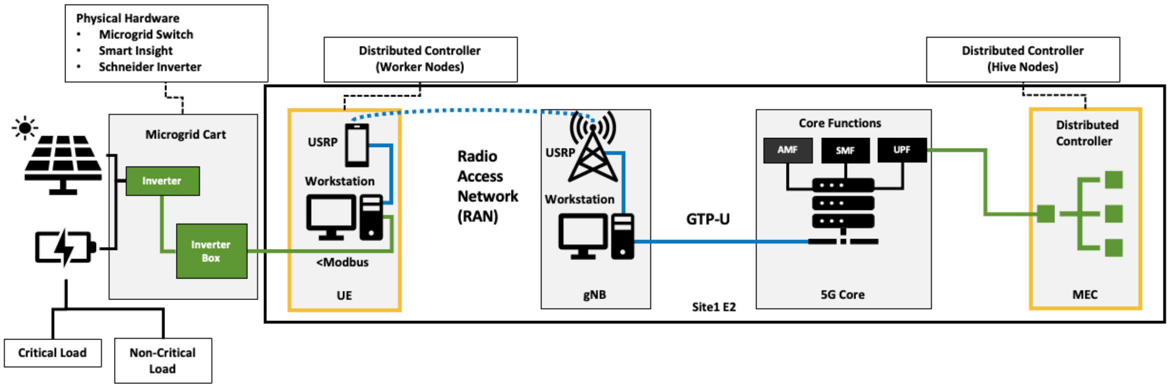

Figure 1. Grid-edge wireless communication for control system

AMF = access management function; gNB = gNodeB; GTP-U = general packet radio service tunneling protocol user plane; MEC = multi-access edge computing; SMF = session management function; USRP = universal software radio peripheral; UE = user equipment

[2] https://www.nrel.gov/docs/fy24osti/87180.pdf

The custom DCS and physical 5G microgrid developed in Task 4 was built and tested at the National Renewable Energy Laboratory's (NREL's) Flatirons Campus near Boulder, Colorado. The Flatirons Campus is a 300-acre field test environment containing a variety of experimental energy resources up to 20 megawatts in scale. For this project, the Flatirons Campus provided an open range for 5G experimentation with geographically dispersed energy resources. There are metal buildings at each site, some of which contain 5G network core racks and some of which contain power grid/microgrid and energy controllers. Figure 2 shows the wireless communications grid control field sites on the Flatirons Campus for Task 4. Sites 1E.1, 1E.2, and 3.3 were used as field sites to perform all PHIL test cases and controls scenarios. Sites 1E.1 and IE.2 provided the contained MEC, OAI 5G RAN, and local control integration with PHIL components. Site 3.3 housed the Celona 5G RAN and local control integration with PHIL and integrated with the MEC located at Site 1E.2.

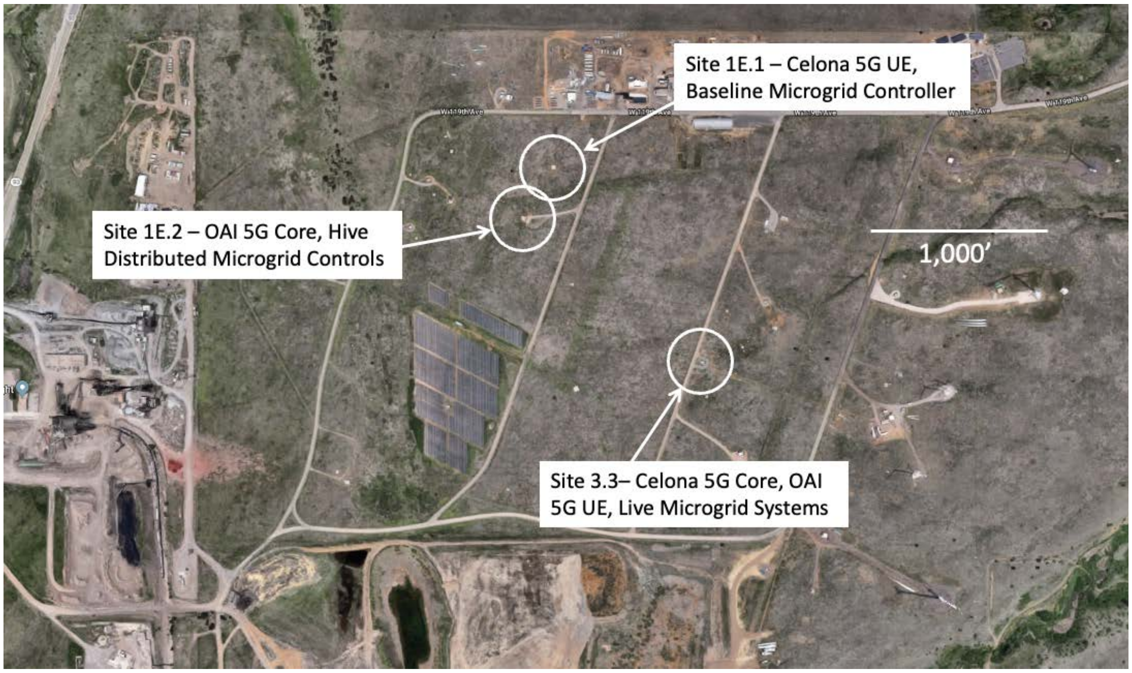
Site 1E.1 – Celona 5G UE, Baseline Microgrid Controller
Site 1E.2 – OAI 5G Core, Hive Distributed Microgrid Controls
1,000'
Site 3.3– Celona 5G Core, OAI 5G UE, Live Microgrid Systems

Figure 2. Wireless communications grid control field sites on NREL's Flatirons Campus

1.1 Distributed Control System

The DCS was developed to manage grid-edge power system devices such as power line circuit breakers, controllable loads, and smart inverters in our testing. The method for controlling these devices differs from traditional centralized control in that there is no master controller gathering data and dispatching control to edge devices. Instead, a swarm of lightweight controllers collaborate to share relevant data and influence each other's control of specific edge devices. One benefit of a centralized approach is the omniscience of collecting all relevant data into a single point. In a decentralized system, timeliness of data may be sacrificed. However, the trade-off of omniscience for modularity, resilience, segmentation, and dispersed computing with our decentralized approach has not compromised our ability to manage a collection of grid-edge devices. For these reasons we have developed our DCS as a prototype/proof of concept for an innovative next-generation control scheme when paired with 5G infrastructure via the MEC and UE.

The DCS in Figure 3 is a network consisting of “hive nodes” (servers) and “worker nodes” (clients) collaborating to provide resilient automated control of interconnected power systems. The DCS contains a mesh network of hive nodes designed to disseminate information throughout the distributed network and corresponding worker nodes designed to provide fast, responsive control to the power systems at the grid edge. The details on how the DCS achieves these goals are described in further subsections.

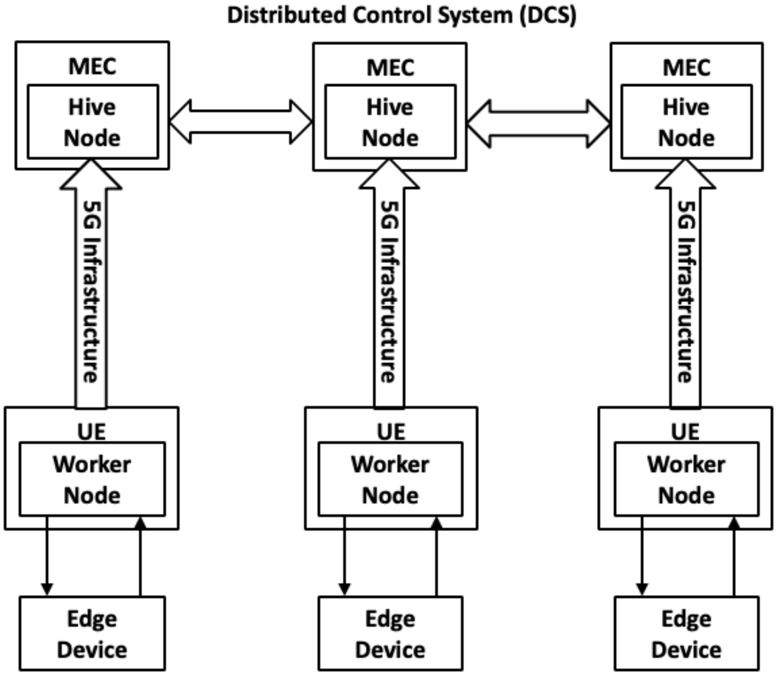

Figure 3. Distributed control system

1.1.1 DCS Requirements

The design of our distribution controller was heavily influenced by the unique requirements of this project. To have timely control of grid-edge devices, latency must be minimized. The acceptable latency depends on the use case, system specifications, communication protocol specifications, and any other standards the whole system is being compared to. For this project, we aimed to have the DCS decision and response functionality below the common polling rate of Modbus protocol, about once every 100 milliseconds (ms). The 5G and microgrid interface design decisions will impact the overall system latency. Table 1 lists the latency compounding design challenges considered while developing the DCS.

Table 1. DCS Latency Compounding Design Challenges

Check Point #	Latency Compound Design Checks
Latency Check Point 1	For the situation where one edge device changes state and the local controller needs to react to that state change by controlling a separate device, there are eight places where latency will compound. In the chain of actions between a device changing state to the reaction from that changed state, latency is compounded at each step. The latency is not constant but varies between each step; nonetheless, each step will add some measure of latency.
Latency Check Point 2	In the situation where the distributed control system needs to react to the changing state of grid-edge devices across neighbors, there is extra processing and a sharing action resulting in 10 steps where latency is compounded. There is an added step of sharing the state change information with a neighboring node. This step comparatively adds a large amount of latency to the final calculation. In our project we deliberately imposed this state to simulate the latency of a long-distance communication by adding 200–250 ms of latency. The ~200 ms latency was chosen as a realistic value based on existing research on latency between regional cloud data centers. (Chen et al. 2021). The goal of adding this latency is to simulate a realistic scenario of geographically dispersed distributed computing. This is needed because in our testing architecture, the distributed controllers are physically in the same location.

The decisions made in designing our DCS focused on reducing latency as much as possible. These design decisions are explained in greater detail in further sections. Other system requirements include correctness of control logic interpretation, modularity, low coupling with hardware-in-the-loop, and synchronous and asynchronous capabilities. These requirements were validated through scenario testing and internal code review.

1.1.2 Distributed Control System Architecture

The DCS is designed in a modular fashion with hive nodes and worker nodes, as illustrated in Figure 4. The hive node is programmed using Golang, and the worker node is programmed in C. These languages were chosen because they are fast and efficient, which is a requirement to reduce processing latency. We designed the hive node using Golang to take advantage of the simple and performant multithreading capabilities of the language. We used C for the worker node for its performance and because the manufacturing message specification (MMS) and Modbus protocol libraries used to interface with edge devices were also written in C.

The details on how each hive node's modules function in relation to each other and to worker node threads are expressed in the following sections. For more information about software versioning, libraries, and specifications, see Table A-4 in the Appendix.

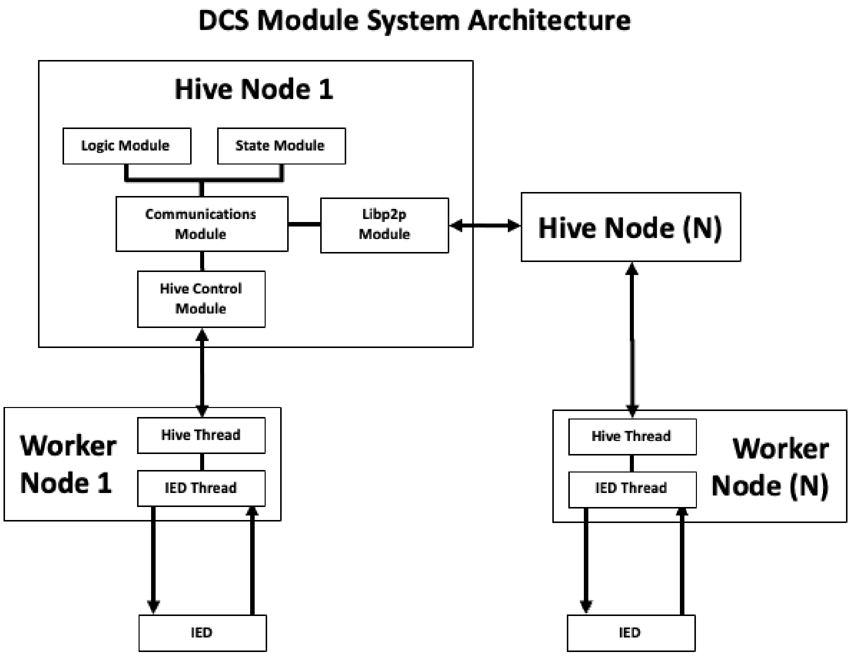

Figure 4. DCS module system architecture

1.1.3 Hive Nodes

The hive nodes of the DCS comprise a series of software modules that establish the overall functionality of the hive node system architecture. These hive modules make up the overall functionality for communications, controls, data sharing, logic, and state monitoring across the hive system using distributed peer-to-peer communications. MMS and Modbus are used by worker nodes to manage the intelligent electronic device (IED) grid edge. Table 2 provides a more in-depth description of each of these modules and their functionality in the context of the DCS.

Table 2. DCS Modules

Communications Module	The communications module defines and manages the input and output channels used for each module to communicate with each other. There are four channels: input from neighboring hive nodes, output to neighboring hive nodes, input from registered worker nodes, output to registered worker nodes.
Hive Control Module	The hive control module manages connections to registered worker nodes. This module receives state updates from the worker nodes and sends policy updates to worker nodes. This module pulls messages from the worker output channel and pushes messages to the worker input channel.
Libp2p Module	The Libp2p module manages connections to each neighboring hive node. The main goal of this module is to process incoming and outgoing messages for neighboring nodes. Two threads are created, one for receiving messages and inserting them into the hive input channel and one for retrieving messages from the hive output channel and broadcasting them to neighboring nodes.
Logic Module	The logic module defines and manages the control policies used by worker nodes to directly control grid edge devices such as circuit breakers, controllable loads, and inverters.
State Module	The state module defines and manages a representation of the shared state held between all hive nodes in the distributed system. Each hive node creates an initial state, which is shared with the neighboring hive nodes. As the worker nodes update their respective hive nodes, the state representation is updated and shared to neighboring hive nodes. This state is shared between all nodes in the distributed network until all nodes are synchronized.

1.1.4 Worker Nodes

The worker nodes are constructed of two threads, one for handling interactions with the hive node and one for handling interactions with the IED. These threads make up the overall functionality of the work in its relation to the hives. Much like the hive nodes, information needs to be passed between the two threads to relay IED status updates back to the hive network. Table 3 provides a more in-depth description of each thread and its functionality in the context of the DCS. Figure 4 shows the relationship between worker and hive nodes.

Table 3. DCS Threads

Hive Connection Thread	The hive thread is responsible for connecting, receiving from, and sending to the hive node registered with the worker. Incoming information encompasses registration messages and policy update messages. Outgoing information encompasses state updates from the configured IED.
IED Connection Thread	The IED thread is responsible for connecting, receiving from, and sending to the configured IED. Outgoing information encompasses direct control commands using MMS or other industrial control system control protocol. Incoming information encompasses solicited and unsolicited IED status updates.

1.1.5 DCS Communication Relationship

The DCS communication model has gone through several iterations over the course of this project. The original design consisted of a single hive node directly connected to neighboring hive nodes and an assigned edge device. Throughout the project, we gained insight into the incompatibility between the 5G control paradigm and the grid-edge control paradigm that we were working with. Details about this incompatibility are further described in section 1.1.6. Our final iteration: hive-to-hive, hive-to-worker, and worker-to-IED is highlighted in Table 4. This final design addresses the requirements stated in paragraph 1.1.1 and is critical to the system's overall purpose to scale as well as maintain a resilient and reliable DCS architecture in the event of communications loss across the MEC, RAN, or grid edge. This fully decentralized DCS approach also features inherent cybersecurity advantages. The primary advantage is due to the limited scope of each distributed controller. Instead of relying on a single point of failure, a DCS node that crashes or that is attacked only affects one device or a small quantity of devices. This is unlike a centralized controller, which manages many devices.

Table 4. DCS Communications Links

Hive-to-Hive	Hives form peer-to-peer connections with reachable nodes using the libp2p library. This library is also the basis for the Interplanetary File System and Ethereum blockchain. These connections are encrypted using Transport Layer Security 1.3. Over this technology, the hive nodes use a pub-sub gossip protocol, "gossipsub," which allows information to be disseminated through a partial mesh network.
Hive-to-Worker	Each hive node is coupled with a single worker node through the 5G infrastructure we built for this project. A transmission control protocol connection over 5G sends Java Script Object Notation (JSON) data representing control policies (an example control policy is in Appendix A Figure A-1). The control policies provide power system set point instruction for how the worker node will control edge-level IEDs.
Worker-to-IED	Worker nodes create a peer-to-peer connection with one of three edge devices from Figure 3 using the MMS or Modbus protocol. Direct device control and retrieval of device status is performed over this connection.

1.1.6 Distributed Energy Resource Control Model

To establish automated grid-edge control, the distributed controller must pull updates from edge devices and push controls to edge devices. Traditionally, in an industrial control system context, a controller establishes connection to an edge device to set up a session where both endpoints can send and receive messages. However, as shown in Figure 5 in a 5G context, the edge device—usually a cell phone—establishes a connection through the telecommunication network to data servers to accomplish an application's desired functionality. This creates a conflict of control paradigms because the interface designs of industrial control system hardware and 5G hardware are fundamentally different.

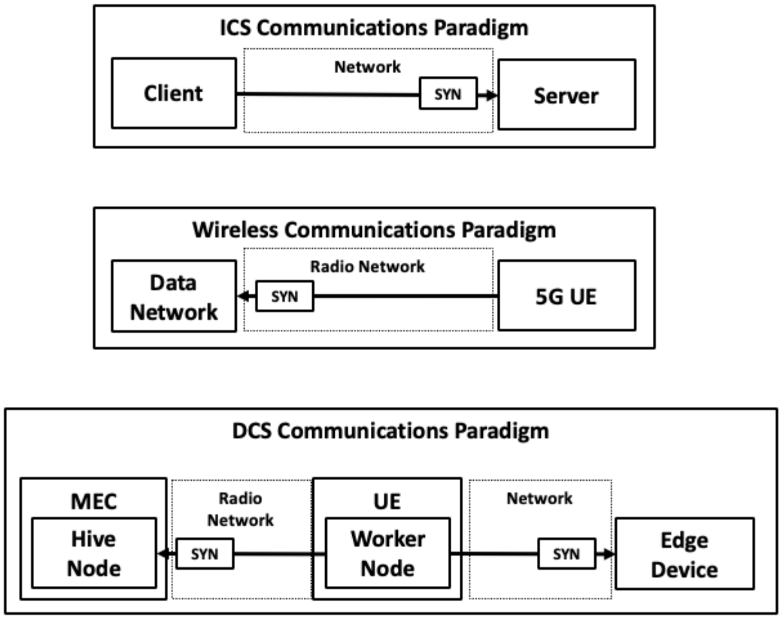

Figure 5. Differences in industrial control system and 5G communication paradigms

In more detail, the primary issue we faced is that the 5G infrastructure requires the edge device to establish the communication session. However, current grid-edge devices do not establish communication sessions due to standards that define the roles and relationships of controlling and controlled devices. Servers (edge devices) listen for client (controller) connections and receive requests. The server will not reach out to a client to establish communication.

With knowledge of this conflict, we designed the distributed controller in two pieces: a controlling (hive) node on the network side of the 5G infrastructure and a proxy (worker) on the device side of the 5G infrastructure. This worker establishes a connection to the device through wired means and a connection to the hive node through 5G. After these communication pipelines are established, the industrial control system control paradigm is possible.

Though the paradigm differs from a security perspective, the inability for communications to come in through the user plane is viewed as a security control. If the user plane (or data network) allowed for communications traffic to pass through the 5G core, through the RAN, to the edge UE by default, attack scenarios against the UE would be more likely to occur. However, since the UE worker node establishes the trust connection through the 5G RAN and user plane before the hive node can provide policy, the threat to the UE work node is reduced, which benefits the overall security and resilience of the communications and controls approach.

The advantages of the constructed worker-hive architecture include:

- Distributing responsibility between hive/worker nodes to increase resilience of edge control
- Bump in the wire solution covers legacy devices
- Worker can be lightweight for fast local reaction
- Maintaining 5G edge paradigm for security and resilience.

1.2 Power Hardware-in-the-Loop

To validate the 5G testbed with physical components, we used microgrid technology located at NREL's Flatirons Campus Site 3.3. This PHIL edge-level system comprises the following components:

Inverter Cart Components:

• Grid-Tied Power Line	NEMA 240v 30a
• Inverter ComBox (Relay)	Schneider InsightHome
• Grid-Forming (Battery) Inverter	Schneider Conext XW
• Energy Storage Battery	Universal UB121100
• Grid-Following (PV) Inverter	Fronius Primo 6
• Load Bank (Noncritical Load)	PTC 650–1,300 W
• 5G Core Rack (Critical Load)	Celona et al.

The microgrid cart (Figure 6) provided the PHIL needed to test the distributed controller and its ability to manage both critical and noncritical loads.

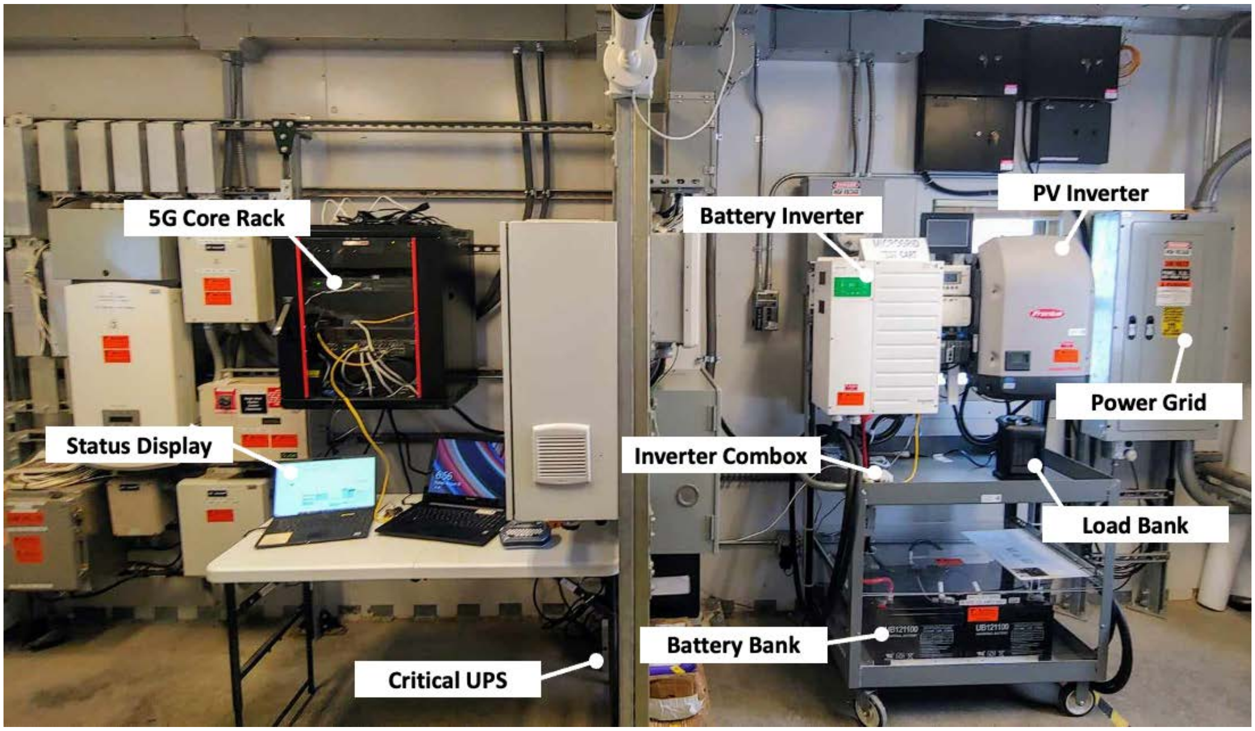

Figure 6. PHIL platform managed by distributed controller (the experiments conducted in this report did not include the Fronius PV inverter)

2 Microgrid Distributed Control System and 5G Radio Access Network Testing

Testing and evaluating the capabilities of the DCS enabled by 5G communication infrastructure was a key objective of this applied research. Our goal was to understand if the combined controller and communications architecture would reliably provide resilient microgrid operations with and without 5G communications. In addition, we wanted to understand how 5G RAN would perform during normal and degraded operations. Operations could be degraded by communications network issues, such as maximum network stress, or by power grid issues, such as a long-duration power outage affecting communications equipment. We conducted test cases to measure system reliability under several scenarios, as described in the following sections.

2.1 System Network Latency

The integration of the distributed controller within the 5G communication architecture for DER control was tested end-to-end for latency. Our latency tests were performed to determine if 5G latency is low enough to enable power restoration within 8 ms (Honrubia-Escribano et al. 2012). If within these boundaries, this could provide uninterruptable "seamless" power restoration. We broke up latency testing into several parts, as demonstrated in Table 5 below. The results of each latency test are presented in Table 5. The anticipated 5G latency was 1 ms or less, but hardware test results only achieved 10–20 ms latency, even after 5G traffic prioritization and optimization. Upon deeper literature research, the "sub-millisecond" latency often quoted for 5G is in reference only to millimeter-wave frequencies (known as "5G frequency range 2" from 28 to 60 gigahertz [GHz]). The frequencies used in this project are "frequency range 1" (FR1) also known as "sub-6 GHz." The 3GPP standard radio time block for FR1 is 1 ms, so it is not surprising that overall end-to-end network latency is higher than 1 ms.

Table 5. System Network Latency Data

Software	Devices	Network	Measured Range	Measured Average	Purpose
Both 5G Solutions	MEC to user plane function	MEC	0.01–0.03 ms	0.02 ms	Validate latency between MEC and user plane function
OAI 5G v1.5.1–v2.0.0	UE to gNB	RAN	5–14 ms	10 ms	Validate latency between UE and gNB
OAI 5G v1.5.1–v2.0.0	UE to MEC	RAN + MEC	6–15 ms	12 ms	Validate latency between UE and MEC
Celona 5G	UE to gNB	RAN	24–34 ms	24 ms	Validate latency between UE and gNB
Celona 5G	UE to MEC	RAN + MEC	26–38 ms	26 ms	Validate latency between UE and MEC
Both 5G Solutions	UE to DER	Grid Edge	0.01–0.04 ms	0.03 ms	Validate latency between UE and DER
Both 5G Solutions	MEC to MEC	Emulated WAN	200–250 ms	n/a	Simulate cross region latency for each message between hive nodes

Based on the observed latencies not meeting the 8 ms threshold, our DCS was re-designed to account for this new understanding as described in section 1.1.2.

2.2 A/B Testing and Controls Scenarios

DCS testing for edge-level DERs was established as a set of scenarios referred to as A/B testing. For Task 4 testing we intended to discover if the system could reliably operate without and with 5G distributed control of the microgrid. Table 6 describes the A/B test cases.

Table 6. A/B Test Descriptions

Test	Description
A	Microgrid operation **without** the distributed controller managing the system.
B	Microgrid operation **with** the distributed controller managing the system.

In addition, a set of controls scenarios were considered and applied to the A/B test cases. These scenarios allowed for a clear representation of what grid operation functionality would look like for a microgrid. These scenarios were chosen to represent the primary functions of a microgrid. A microgrid must transition from grid-following (grid-tied) mode to grid-forming (islanded) mode. This "islanding" can be done deliberately, such as to avoid grid charges or for testing and demonstration, or it can be unplanned, such as when the grid experiences an outage. While islanded, the local energy source is limited and exceeding the limit would cause a full loss of the microgrid power, including the critical load. Therefore, all three scenarios listed in Table 7 are focused on controlling the noncritical load. Scenario 1 turns off the noncritical load when the microgrid's battery gets low, Scenario 2 turns off the noncritical load whenever the microgrid is islanded, and Scenario 3 returns the microgrid to grid-following (grid-tied) mode if the noncritical load is too high.

Table 7. A/B Controls Scenarios

Controls Scenario	Controls Scenario Description
Baseline (Scenario 0)	Human-controlled (manually operated) microgrid.
Scenario 1	Switch off noncritical load at a set battery threshold of 50% discharged. Switch noncritical load on at a set battery threshold of 100% charged.
Scenario 2	Detect when battery inverter has gone into grid-forming mode, disconnect noncritical controllable load.
Scenario 3	When in grid-forming mode (deliberately islanded even though the grid is available), if total load exceeds a threshold, switch to grid-following mode.

We compared the behaviors captured in each test based on the scenario procedures defined in Table 8.

Table 8. Scenario Procedures

Scenario #	Step 1	Step 2	Step 3	Step 4	Step 5	Step 6
Scenario 1	Battery voltage is 57,400 mV or higher (100%)	Wait until voltage drops below 48,300 mV (50%)	Observe distributed controller turn load off	Wait until inverter reconnects to the grid	Recharge battery to 57,400 mV or higher (100%)	n/a
Scenario 2	Start with inverter in grid-following mode	Start with noncritical load connected	Change battery inverter to grid-forming mode	Turn noncritical load off	n/a	n/a
Scenario 3	Start with inverter in grid-following mode	Start with noncritical load disconnected	Change battery inverter to grid-forming mode	Turn noncritical load on	Observe distributed controller set inverter to grid-following mode	Observe inverter reconnect to the grid

The primary goal of the three scenarios is to test whether the DCS can collect, disseminate, and react to data gathered within a microgrid while segregating the responsibility and authority of managing the microgrid. For this scenario, our infrastructure setup is as described in Section 1.2. Per our policy for this test, one worker/hive pair is responsible for collecting data from the inverter. A second worker/hive pair is responsible for controlling the noncritical load as a reaction to state update information. And a third worker/hive pair is responsible for managing the grid-connectedness of the inverter.

2.2.1 A-Testing and Controls Scenarios

The A-Test, which represents a microgrid without the distributed control (or manual operations), used the scenarios listed in Table 9.

Table 9. A-Test Scenario Results

A-Test Scenarios	A-Test Results
Baseline	The baseline is a human-controlled microgrid. The power grid experienced an outage and the microgrid is not automated, so the noncritical load lost power. The critical load has its own uninterruptable power supply (UPS), so it remains powered on but can only last a couple hours. As the power outage persisted, the human operator started a long, multistep process to black-start the microgrid. At the 10-minute mark, the microgrid was operational and was able to provide power (at least for critical loads) for as long as PV generation was sufficient.
Scenario 1	Figure 7 shows that when the grid connection is lost, the microgrid runs off the battery for a time. The battery's voltage slowly decays before a voltage spike occurs when the noncritical load is disconnected. When the grid is restored, the voltage rises again as the battery is recharging.
Scenario 2	The inverter was deliberately placed into grid-forming mode. The manual (human) operator then successfully shut off the noncritical load, which conserved energy while the microgrid was islanded, ensuring maximum battery duration for the critical load.
Scenario 3	Figure 7 shows the battery's voltage fall when the grid is disconnected and the microgrid moves into grid-forming mode. Then, when the noncritical load is connected, the voltage falls even farther. After this, an attempt to reconnect to the grid is observed, and voltage climbs again.

Scenario 1

Figure 7 demonstrates grid operations without DCS control (A-Test). The figure represents a time series of voltage data visualizing when the grid loss took place during Scenario 1 (defined in Table 9), when the load turned off, and then when the grid was restored and the load turned on. These data represent what we would expect from the power system in response to A-Test Scenario 1.

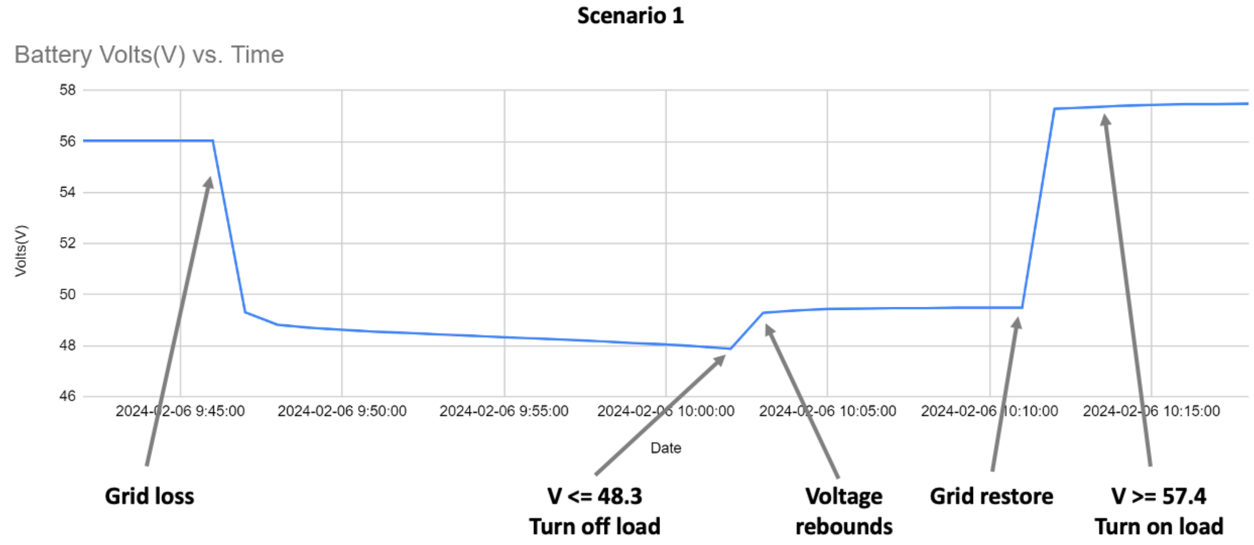

Figure 7. A-Test output for Scenario 1

Baseline / Scenarios 2 and 3

Figure 8 demonstrates grid operations without direct control (A-Test). The figure represents time series data based on the series of A-Test scenarios defined in Table 9. Voltage data from the dashboard, also known as a human-machine interface (HMI), visualizes how the grid behaved during the baseline, Scenario 1, and Scenario 2. The dashboard graph data capture from the inverter HMI represents what we would expect from the defined scenarios.

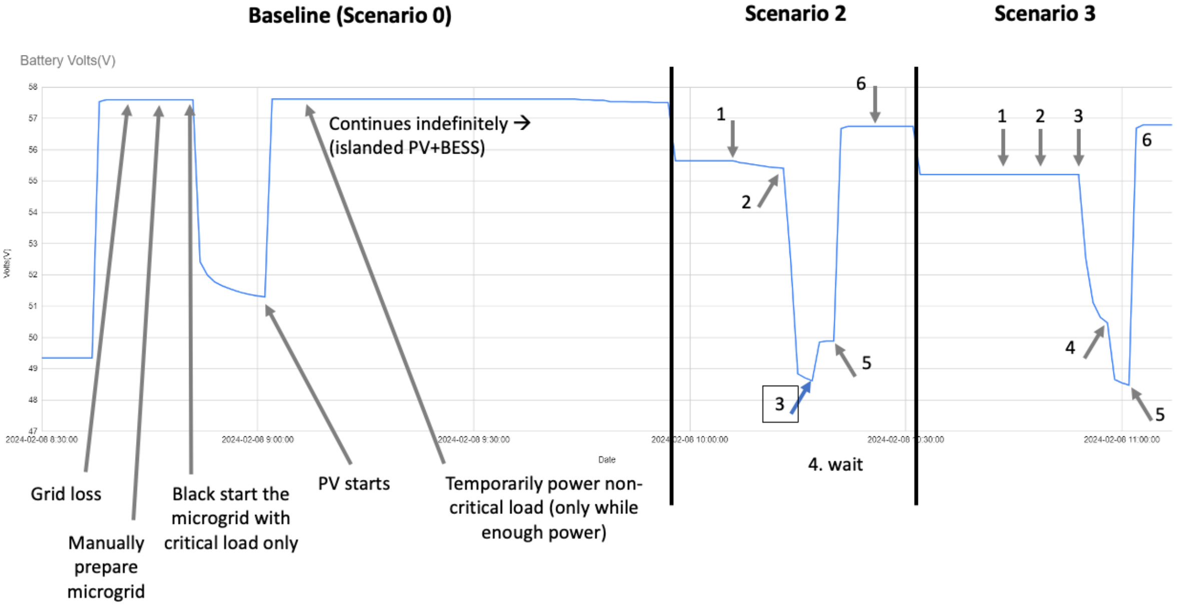

Figure 8. A-Test results for Scenarios 2 and 3 (legend for steps 1–6 found in Table 8)

2.2.2 B-Testing and Controls Scenarios

The B-Test represents a microgrid with the distributed controller. The B-Test with the DCS performed the commands to the grid in a time series that was too rapid to capture any meaningful dashboard visualization from the microgrid-edge HMI. However, Figure 9–Figure 17 demonstrate when the commands were issued via the DCS. Table 10 describes the B-Test scenarios and results.

Table 10. B-Test Scenario Results

B-Test Scenarios	B-Test Results
Scenario 1	The power grid experienced an outage, and the microgrid went onto battery power. The distributed controller observed the battery voltage dropping, and at the predetermined level (50% state of charge), the distributed controller successfully shut off the noncritical load. This conserved energy until the power grid was restored, ensuring maximum battery duration for the critical load. Once the grid was restored, the distributed controller observed the battery voltage rising (charging), and at the predetermined lever (100% state of charge), the distributed controller successfully turned on the noncritical load.
Scenario 2	Scenario 2 turns off the non-critical load whenever the microgrid is islanded, such as when the power grid is down. The distributed controller detects when the inverter has gone into grid-forming mode and responds by disconnecting non-critical controllable load. This conserved energy while the microgrid was islanded, ensuring maximum battery duration for the critical load. Scenario 2 was integrated with Scenario 1 during the testing of the distributed controller; therefore, no data artifacts were collected separately.
Scenario 3	The inverter was deliberately placed into grid-forming mode. The distributed controller successfully detected that the inverter was islanded and observed the electrical load amount. The load was greater than the threshold, so the distributed controller successfully shut off the noncritical load, which conserved energy while the microgrid was islanded, ensuring maximum battery duration for the critical load.

Due to the very rapid nature of the 5G distributed controller, the HMI for the PV inverter was not able to establish a graph representative of state changes (all changes took place within one time step of the graph). Therefore, packet capture data gathered directly from the 5G UE worker-to-IED interface allowed us to establish a series of graphs representing a time series of state data (voltage/load) for the power system and grid. We were able to clearly see when the control systems applied actions and performed logic across the system. The following set of artifacts helped us clearly see that the distributed control system successfully performed Scenarios 1 and 3.

Scenarios 1 and 2

For Scenario 1, we began our test in the initial state with the microgrid in grid-tied mode, the noncritical load active, and the critical load active. With the DCS active, the observed state of the inverter, battery voltage, and total load measurements were continuously pulled and relayed to the distributed network of hive nodes.

After recording a stable state, the inverter was manually transitioned to islanded mode through the HMI to simulate a grid outage. Shown in Figure 9, Figure 10, and Figure 11, this event occurs near the 150-second mark of the test. We observed the battery voltage drop steadily until it reached the Scenario 1 policy low-voltage threshold of 48.3V near the 650-second mark in the test. At this moment, the DCS disconnected the noncritical load to preserve battery voltage for as long as possible. After observing a steady state for a short period of time, we manually reconnected the grid through the inverter HMI.

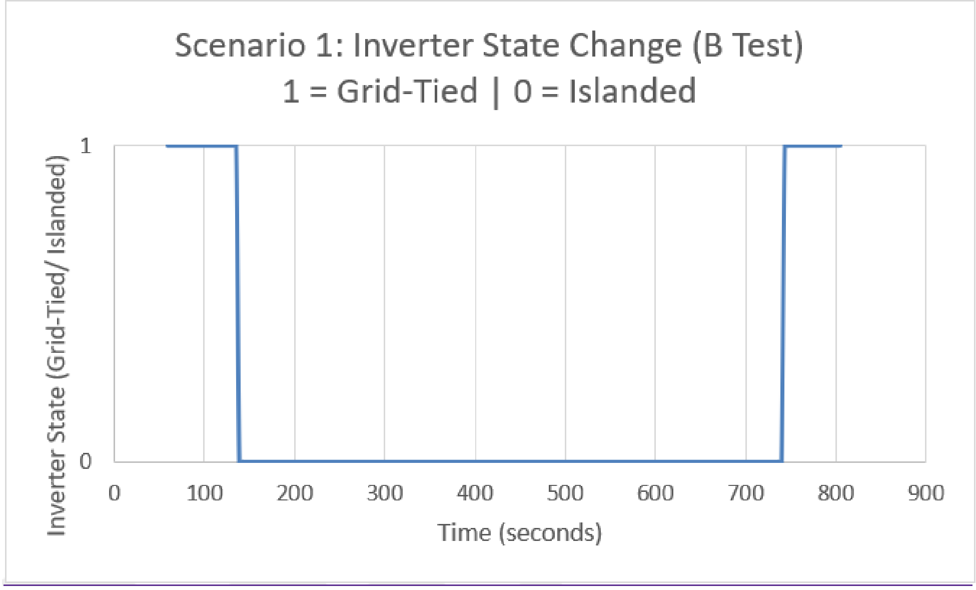

Figure 9. Scenario 1 inverter state change over time

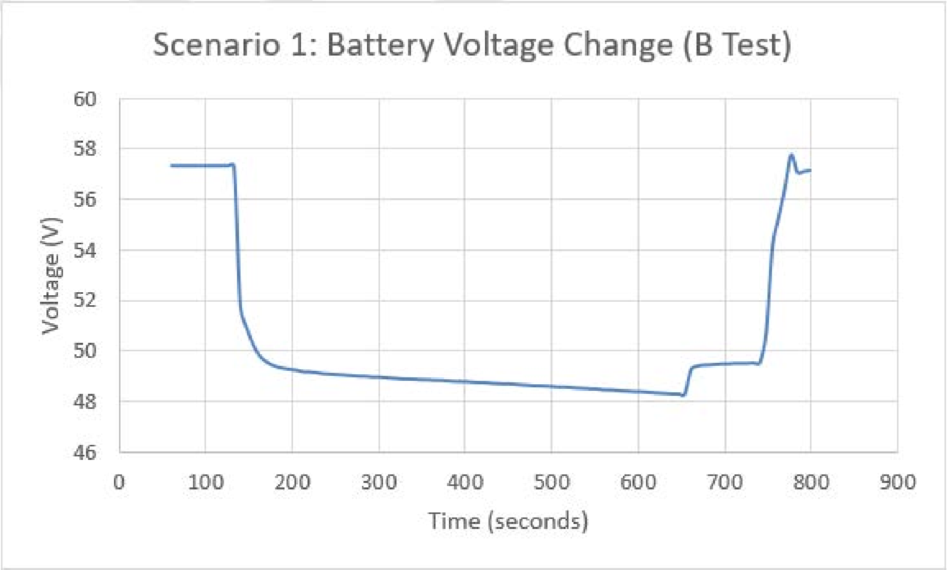

Figure 10. Scenario 1 battery voltage change over time

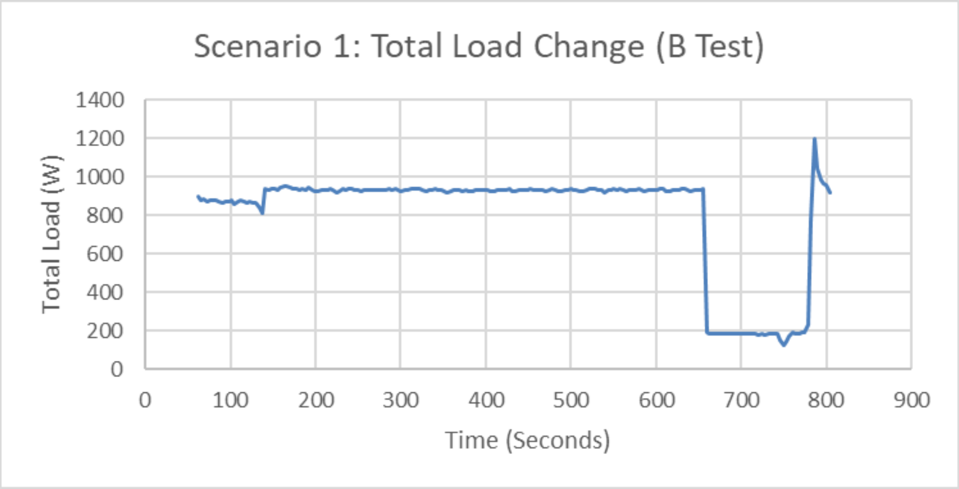

Figure 11. Scenario 1 total load change over time

The grid reconnection caused the battery voltage to increase from the 750-second to the 800-second mark of the test as it recharged. Finally, when the battery voltage reached the Scenario 1 policy high-voltage threshold of 57.4 V the DCS commanded the noncritical load to turn back on, restoring normal operation.

After testing Scenario 1, we determined Scenario 2 was subsumed by Scenario 1. Our goal with this test was to prove the reactive functionality of the DCS given a dynamic environment. The more complex test in Scenario 1 shows the DCS is more than capable of the simple state change of Scenario 2. In the interest of time for this project we decided to forego testing Scenario 2.

Scenario 3

Figure 12 demonstrates the state change of the inverter based on the actions defined by B-Test Scenario 3 in Table 10. During Scenario 3's 100-second time series presented in the graph, the data highlights when the inverter state shifted from grid tied to islanded and then back to grid tied between 40 and 60 seconds. This graph represents what we would expect the inverter to do in relation to the distributed control system's automated response.

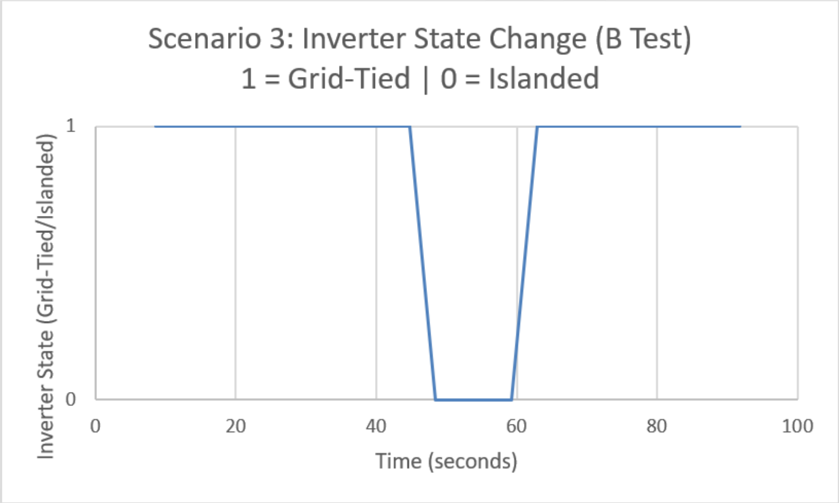

Figure 12. Scenario 3 inverter state change over time

Figure 13 represents the battery voltage of the inverter based on the actions defined by B-Test Scenario 3 in Table 10. During scenario 3's 100-second time series presented in the graph, battery voltage dropped from just over 56.68 V to 51.3 V and recovered to 56.69 V in lockstep with the inverter in Figure 12.

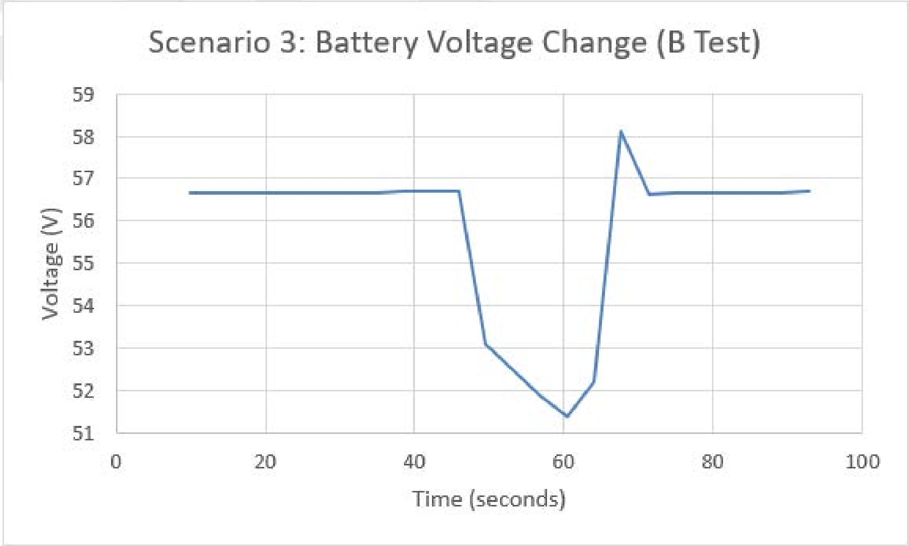

Figure 13. Scenario 3 battery voltage change over time

Figure 14 represents total load change on the actions defined by B-Test Scenario 3 in Table 10. The data represents the 100-second time series when the load change shifted.

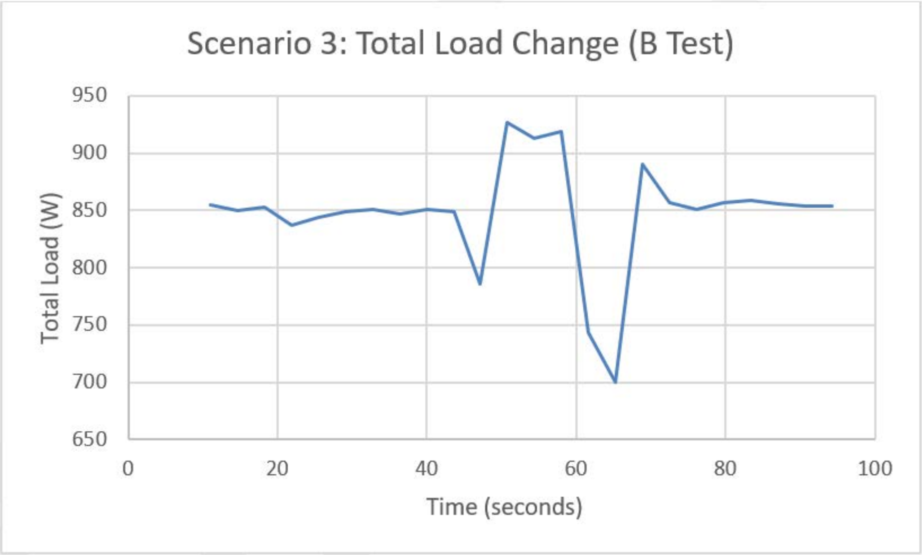

Figure 14. Scenario 3 total load change over time

2.3 Communications On/Off Testing

To test the resilience of the distributed controller and its integration with the microgrid hardware, we performed communications (or "comms") on/off testing to discover how the grid would operate with and without communications to the microgrid.

The comms on/off test compares two tests:

- **On**: the microgrid distributed controller is connected and able to react to the system
- **Off**: microgrid operation without any networking connected

The comms "on" tests are shown earlier—for example, the A/B test scenario results, which relied on comms being on. Scenario 1 was repeated with comms "off" and the results were observed. This scenario begins with a long-duration power outage. Even without comms, the microgrid islanded itself and maintained power to all loads. However, the battery soon dropped below 50% charge and, because of no comms, the 5G control did not turn off the noncritical load. Thus, the battery direct current (DC) voltage continued to fall quickly. At the 18-minute mark, the inverter displayed a "DC under voltage" warning, and 1 minute later the inverter went offline due to battery depletion. Without the inverter, there was no power for noncritical load; thus, noncritical load went offline. The critical load has its own UPS battery, so it remained powered for an additional 41 minutes. After this time, the UPS battery was depleted, and the critical load (the 5G core network) went offline. Table 11 describes the comms on/off test cases.

Table 11. Comms On/Off Test Cases

Test Case	Hive-to-worker	Worker-to-IED	Power Grid	Result
Comms On	Hives provided control policies to 5G UE workers	Policy was applied to 5G UE worker and grid edge	System maintained reliable grid communications	All scenarios succeeded in conserving and restoring electrical power
Comms Off	Hive nodes unable to connect and lost link to distributed control worker	N/A - no link	Systems continued status quo until batteries depleted	Eventual power loss on critical load because the microgrid controller was unable to conserve power; UPS battery failed after 1 hour

2.4 Controls Stress Testing

We performed stress testing on the communication and power layers to discover the resilience and reliability of the 5G RAN, DCS, and overall grid. We expected the 5G RAN to be stable and the DCS to maintain reliable control of the power grid, even when communications channels were stressed. Table 12 lists the stress tests and results.

Table 12. Stress Testing Categories

Stress Testing	Description	Result
Communication Stress	Apply bandwidth consumption across each 5G RAN and MEC to identify the ramifications on the grid controls while operating scenarios. Full bandwidth was consumed by applying a traffic generator such as iperf (used for OAI 5G RAN) and OpenSpeedTest (used for Celona 5G RAN).	The consumption of bandwidth across each 5G system demonstrated no impact to the distributed controller in context of policy sharing and state monitoring. Latency increased across the OAI RAN; however, no direct impact was observed on the energy system. During max bandwidth tests with both OAI and Celona, the controller was still able to operate the power systems. This was expected due to 5G traffic prioritization (5QI). Microgrid controller traffic was given highest priority, and all other traffic, such as stress traffic, was given default priority. However, when 5QI was incorrectly configured (stress traffic given highest priority) on the 5G system, the stress test failed, as demonstrated in Figure 18.
Power Stress Test	Comparing the power draw in idle and stressed state of the RAN using OAI and Celona.	During maximum network stress (full bandwidth traffic), the OAI core network did not draw significantly more electrical power than at idle. However, the fully stressed Celona core network consumed 153% electrical power relative to when network was idle.

Figure 15 demonstrates the state change of the system during communications stress. The data points in the graph map out the state change of the inverter throughout the time series. At the 50-second mark, the inverter state changes from grid tied to islanded. Soon after, the noncritical load is turned on, and we observe a spike in total load. This total load surpasses our threshold for the scenario. As a response, the inverter is commanded to prioritize grid connection. This behavior represents what we would expect as normal behavior for this scenario, showing the stress test did not impact functionality of the DCS.

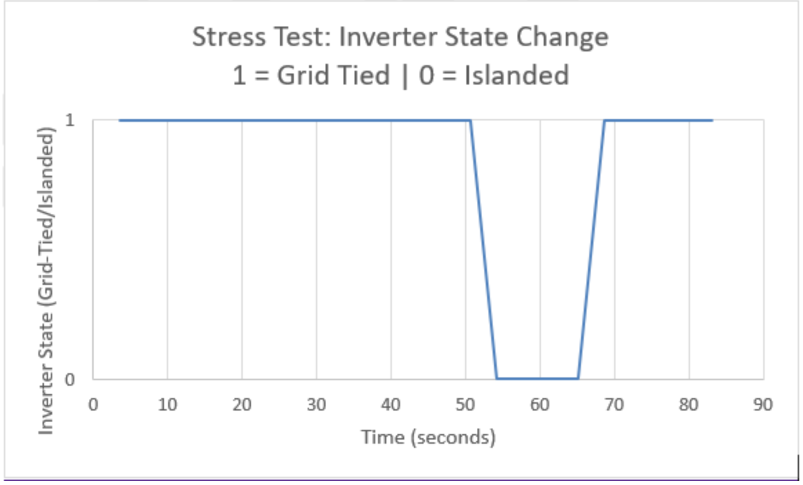

Figure 15. Inverter state change with communications stress test

Figure 16 represents grid tied to islanded back to grid tied during communications stress testing. The graph demonstrates the impact the communications stress had on battery voltage. At 50 seconds the battery voltage changes from 56.8V to 51V. This state change from the inverter is what we would expect to see as the system adjusts to the stress.

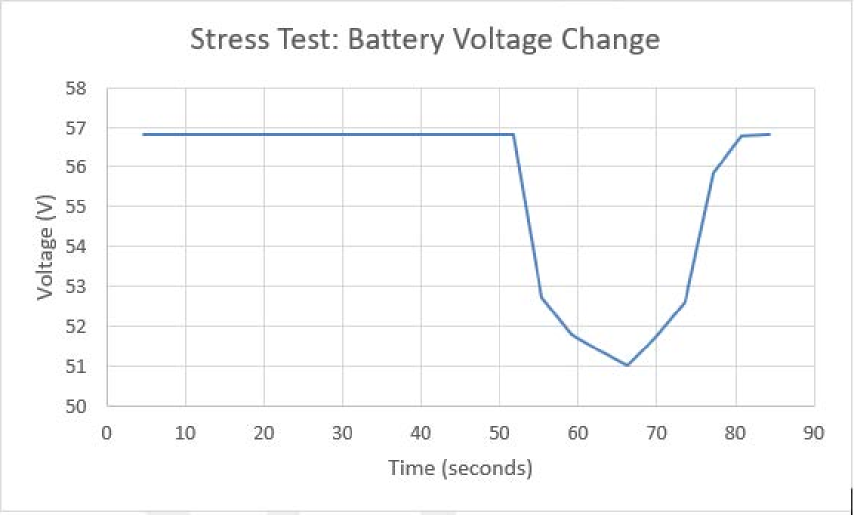

Figure 16. Battery voltage change with communication stress test

Figure 17 shows total load change during communications stress testing. The data illustrates that the total load increases at 50 seconds from 860 W to 960 W, which correlates to previous inverter state and battery voltage change graphs during the communications stress interval. The data point at 50 seconds correlates correctly with the previous graph representing the state change of the inverter.

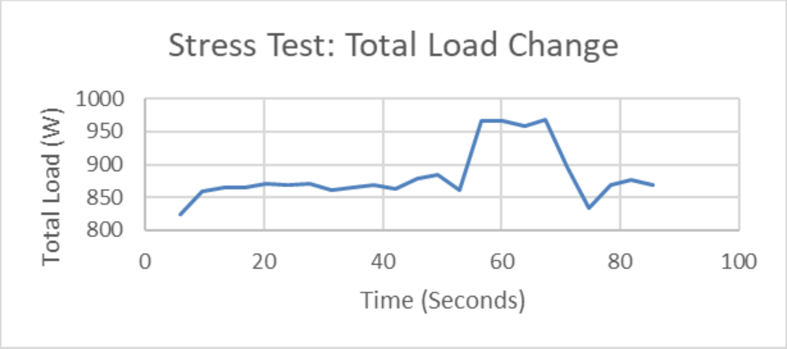

Figure 17. Total load change with communications stress test

Figure 18 demonstrates the outcome of a failed stress test, which resulted from misconfiguring the 5G Quality of Service Identified (5QI). The inset image on the right is a display of the stress traffic bandwidth. During earlier latency testing, the stress traffic (OpenSpeedTest server running at a maximum bandwidth) was set to highest priority, and thus the microgrid controller traffic failed. As shown in top left of Figure 18, the Modbus connection failed, and as shown in the center background, the microgrid dashboard suffered "connection lost." It was not possible to connect to the microgrid, view microgrid data, or control the microgrid.

Once the 5QI prioritization was properly configured with microgrid traffic as highest priority and all other traffic (such as stress traffic) as "default" priority, the test was successful. Although the stress traffic still occupied all the bandwidth, the microgrid traffic worked perfectly because it preempted all lower-priority traffic. The connection to the microgrid was immediate: microgrid data were displayed on the live dashboard, and controls were successful on the first attempt.

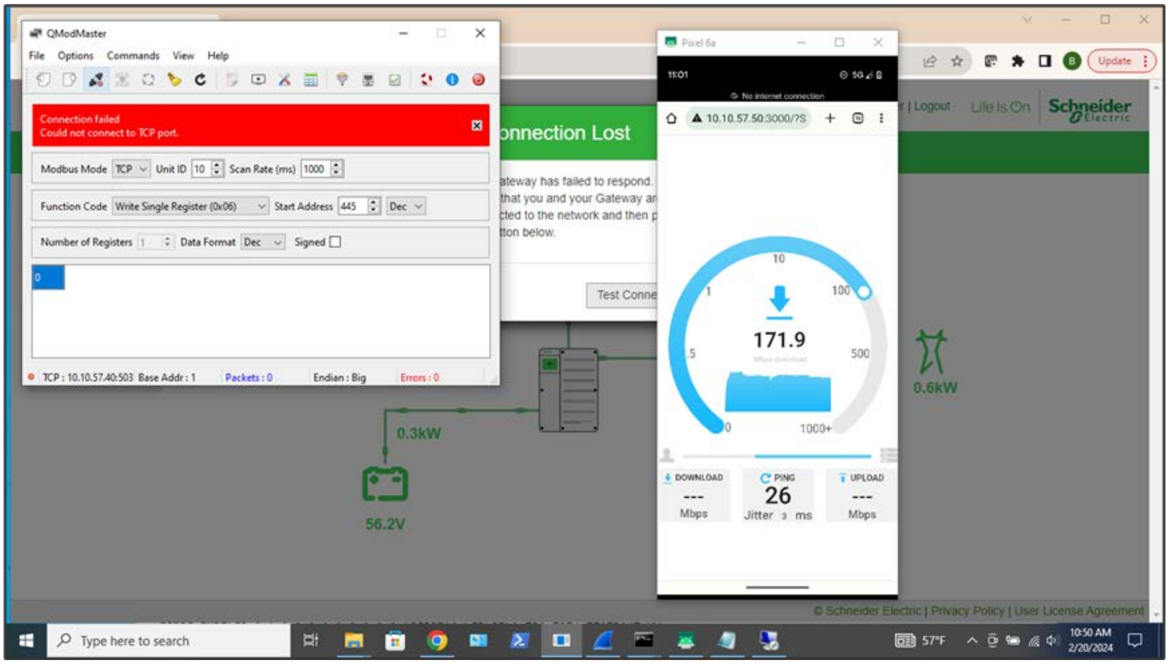

Connection failed
Could not connect to TCP port.

Modbus Mode TCP Unit ID 10 Scan Rate (ms) 1000

Function Code Write Single Register (0x06) Start Address 445 Dec

Number of Registers 1 Data Format Dec Signed □

0

TCP : 10.10.57.40:503 Base Addr : 1 Packets : 0 Endian : Big Errors : 0

onnection Lost

ateway has failed to respond.
that you and your Gateway ar
cted to the network and then p
tton below.

No internet connection

10.10.57.50:3000/?S

DOWNLOAD --- Mbps
PING 26 Jitter 3 ms
UPLOAD --- Mbps

Figure 18. 5QI misconfiguration stress test failure

3 Broader Systemwide 5G Operations

Based on the previous test cases, we considered the implications in the context of broader systemwide 5G operations for a series of edge-level microgrid scenarios. The goal of this theoretical analysis is to understand how the tested systems might perform at scale in a deployed power system.

3.1 Cross-Region Model

We developed a cross-region communications model, which includes multiple 5G network operators (cores), 5G towers, and grid-edge (electrical power grid) equipment devices. It is shown in Figure 19 that electrical equipment may communicate with multiple towers, then the 5G edge applications (distributed control applications for power grids) collocated at each tower's MEC server can communicate with MEC applications at other towers, even if they are on different networks/cores.

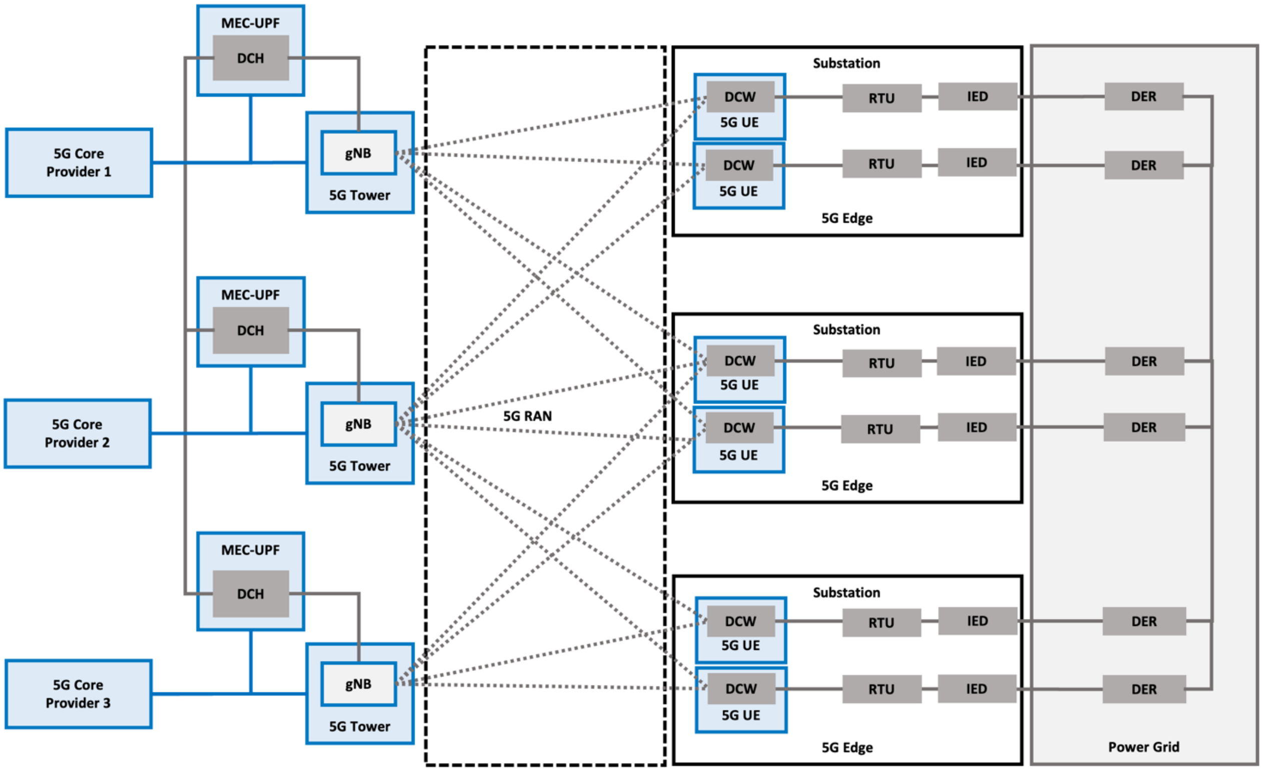

Figure 19. At-scale 5G interconnected DER architecture

DCH = Distributed Control Hive; DCW = Distributed Control Worker; RTU = Remote Terminal Unit

This cross-region model and integration approach for power system operations is critical to enabling a resilient and reliable communications infrastructure for multi-region distributed control scenarios. By providing a fully meshed controls and communications system, operators will be able to ensure that cross-region communications and controls are possible for the energy system during communication outages. In addition, regional systems can inform nearby regions of events to the power grid, enabling situational awareness of weather changes and/or threats.

3.2 Long-Duration Power Outage

We completed a test with the OAI 5G core rack powered by a rackmount UPS battery in order to determine what happens in the event of a long-duration power outage. As expected, without a microgrid, the OAI 5G core network only survived until the UPS battery became depleted. This test was then repeated with a microgrid containing the 5G Celona core, gNB (tower), and MEC server as the critical load. Again, a long-duration power outage was imposed, but this time the microgrid restored power to the UPS battery and, in turn, the critical load. Eventually the microgrid energy storage became depleted (since the PV system was not producing power, all the power came from the battery) and the critical load switched to getting power from its own UPS. For this non-5G-controlled baseline test, the power grid remained down for a long time, which means the critical load's UPS battery eventually ran out, and the 5G core/tower/MEC failed due to power loss.

For the 5G-controlled grid restoration test, the power grid again went down but was restored within seconds (time for the 5G distributed controller workers/hives to restore the power grid). Since the power grid was restored before the microgrid's energy storage ran out, critical power for the Celona 5G core/tower/MEC did not ever fail.

For a broader system, it is assumed that some 5G nodes will never fail because the 5G-enabled distributed controllers restore the power grid quickly after an outage—quickly enough that UPS units or microgrids can successfully endure. However, we also assume some 5G nodes may be located on a damaged power line (at the source of the power grid outage) and thus power cannot be restored quickly, not even with 5G-enabled automation, and will only remain online for 2 hours (assuming a typical 5G node has a 2-hour battery and no backup generator).

Per the scope of work, we analyzed a hypothetical scenario based on Marine Corps Air Station Miramar, with the addition of one 5G node off base (battery backup only) and one on base (full microgrid backup). The objective of the Miramar scenario was to hypothetically extrapolate the test results to see how they might impact a realistic military base.

In the Miramar scenario, the critical military installation experiences a long-duration utility power outage. Because the 5G towers on and off base have 2 hours of battery UPS backup, the communication networks stay online initially. The Miramar microgrid restores power to critical portions of the base, including the on-base 5G tower, long before the tower's UPS battery depletes. However, the off-base 5G tower eventually goes offline once its UPS battery is depleted. Because the coverage area of the towers overlaps, all the network traffic is handed to the on-base tower. Thanks to 5G network prioritization, the Miramar traffic is not impacted, but the off-base traffic may be throttled (except for emergency calls, power system controls, and other prioritized types). Overall, the Miramar 5G communications and electrical power for its critical mission stay online throughout the emergency. The 5G systems and the power systems worked to support each other synergistically as predicted. Even when a loss of 5G communications was forced, the power systems remained powered, but without the longer (or indefinitely long) survivable duration as enabled when noncritical load is intelligently controlled.

4 Summary and Next Steps

The 5G Securely Energized and Resilient project achieved key milestones toward the integration of 5G communications with an energy network to provide secure and resilient operations. Our goal was to evaluate and use 5G features to improve overall microgrid operations by leveraging innovative distributed controls while managing system energy demands to extend operational duration. The project outcomes have significant dual-use impact for the U.S. Department of Defense, other government agencies, and the commercial sector. Our research efforts were organized into several tasks over the past 3 years, and this report is the culmination of the overall project with a focus on the outcomes and insights from the final task. The final task (Task 4) of the project tested the communications infrastructure enhancements and newly developed distributed controls functions interfaced with physical power systems. Earlier tasks included background system research, integrating, and configuring the 5G platform and power system components, and testing with simulated components. The research team successfully completed tests with active power system components in microgrid architecture managed by unique 5G-enabled distributed controls for resilient operations.

We completed testing of several grid operations, communications loading, and system failure scenarios. The latency of the key linkages in the network communications chain were individually measured. We found greater latency on the 5G RAN connection than expected, thus the distributed controller had to be moved closer to the grid edge to achieve satisfactory latency between the worker (controller) and the DER. The impacts of communications bandwidth and loading noted an impact in energy demand on the commercial solutions but less so for the open-source 5G implementation. The reason for the discrepancy is not clear. Future work could further explore the opportunities for energy efficiency improvements to 5G network tech and operating strategies. Finally, we tested several scenarios of the microgrid operations with and without an operational distributed control system deployed in 5G. In each case, we were able to maintain an operational microgrid and communications architecture that was only limited by the available energy resources. In a future at-scale deployment, goals for 5G sustainment along with other critical loads will factor into microgrid scale and controller state decisions.

Enabling the growth of 5G systems, including open radio access area network technology, offers potential value for designing and operating future secure power systems. Our research included developing and maintaining a continuous improvement/continuous deployment architecture that enabled rapid network configuration, monitoring, and reconstruction with confidence that the system architecture would work as intended based on individual software component tests. By using both commercial and open-source solutions for 5G, we captured key priorities for cybersecurity actions. These outcomes were summarized in prior reports and presentations.

Innovations resulting from this work include:

- Developing the 5G edge node integration with power systems components control parameters
- Using a hive-worker with a proxy for data flow and controls execution with managed security
- Controlling the 5G components to conserve energy and enable long-lasting microgrid survivability of the communications and critical component loads

- Developing a robust deployment strategy for 5G technology that will enable cybersecurity strategy implementation.

In addition to documenting our innovations, all the virtualized 5G components have been packaged and ported into the NREL Advanced Research on Integration Energy Systems Cyber Range to enable continued research and development of 5G for energy systems networks.

Future work should include establishing trusted interoperable interfaces between 5G cores such that data can be safely shared, and edge device coordination can take place at a significant scale. Additionally, there are analytic functions that should be further researched, including the network data analytics function and the network exposure function. These can enhance the network configurability and operational insights such that grid components could be tuned to efficiently manage consumption, data sharing, and interconnections. Other security features like zero-trust methods and network slice configuration require further investigation. We see an opportunity to quickly move and analyze situation-relevant data within the MEC as an enabler for smart grid operations in the future.

The integration of 5G with energy systems management presents an opportunity for innovation, as was demonstrated in this project. We successfully tested microgrid controls with 5G communications, developed distributed controls that can work within a secure 5G architecture, and highlighted the balance between comms as a critical load and survivability strategies using a microgrid. Testing confirmed some of the initial expectations while also identifying further research needs. Future work should include steps to enable the low latency between edge devices and controls along with strategies that make use of real time data analytics for more advanced security and energy efficiency actions. Given a parallel growth path for both 5G communications and renewable energy, there will be clear opportunities for the Department of Defense and utilities to implement value-added innovations.

Glossary

Term	Definition
FutureG	The next generation of wireless communications.
Golang	A complied programming language used to develop software.
gNodeB	5G radio unit (or base station) to provide cellular service coverage to user equipment.
Grid-tied	Power system connected to the utility-owned power grid.
Islanded	Power system not connected to the utility-owned power grid. System powered by local power generation as a microgrid.

References

Honrubia-Escribano, Andres, Emilio Gómez-Lázaro, Angel Molina-García, and Juan Alvaro Fuentes. 2012. “Influence of voltage dips on industrial equipment: Analysis and assessment.” *International Journal of Electrical Power & Energy Systems* Volume 41 (Issue 1): Pages 87-95. https://doi.org/10.1016/j.ijepes.2012.03.018.

Chen, Xusheng, Haoze Song, Jianyu Jiang, Chaoyi Ruan, Cheng Li, Sen Wang, Gong Zhang, Reynold Cheng, and Heming Cui. 2021. “Achieving Low Tail-Latency and High Scalability for Serializable Transactions in Edge Computing.” *EuroSys '21: Proceedings of the Sixteenth European Conference on Computer Systems*: Pages 210–227. https://doi.org/10.1145/3447786.3456238.

Rivera, Joshua, Brian Miller, Jordan Peterson, Eric Feth, Paul Snyder, and Tony Markel. 2023. *5G Securely Energized and Resilient: Task 2 and 3 Progress Report*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5R00-87180. https://www.nrel.gov/docs/fy24osti/87180.pdf.

Appendix: Technical Equipment Details and Example Policies/Libraries

Table A-1. 5G System Software for OpenAirInterface (OAI)

5G System	Operation System	Container Software	UHD	OAI
5G Core	Linux Ubuntu 22.04	Docker Version 20.10.21		5G OAI Core version 1.5.1 – 2.0.0
5G gNB Host	Linux Ubuntu 22.04		UHD 4.4	5G OAI gNB Development Version
5G UE Host	Linux Ubuntu 22.04		UHD 4.4	5G OAI UE Development Version

Table A-2. 5G System Firmware for Celona

System Name and IP Address	Application Name and Version	Purpose
Celona access points; 2 DHCP addresses each	AP21-48 with firmware version 2305.ap.0.0-13 or later	Cellular radio service for user 5G devices
Celona edge; 2 DHCP addresses plus 1 static IP per 5G VLAN	Edge Express with firmware version 2305.edge.0.0-35 or later	Cellular network control on premises
Celona managed Ethernet switch; management IP	Netgear GS324TP with firmware version 1.0.0.43 or later	VLAN segmentation boundaries
Edge workstation; static IP address	AMD 5700G barebones computer with latest Windows or Ubuntu OS	DHCP services, edge computing, speed test service host (for 5G)
Celona system orchestrator (CSO); *.celona.io sas.goog spectrum-connect.federatedwireless.com sm-v4-072-d-gtm.pr.go-esim.com	CSO version 2308 or later	Configure and operate cellular 5G core

Table A-3. Universal Software Radio Peripheral (USRP) Software and Hardware with OAI 5G System

Device	Operating System	FPGA Image	FPGA
5G gNB USRP N310	Embedded Linux Alchemy-Zeus 2021.04	USRP N310 FPGA HG	Xilinx Zynq
5G UE USRP N310	Embedded Linux Alchemy-Zeus 2021.04	USRP N310 FPGA HG	Xilinx Zynq

A.1 Example Worker Node Policy

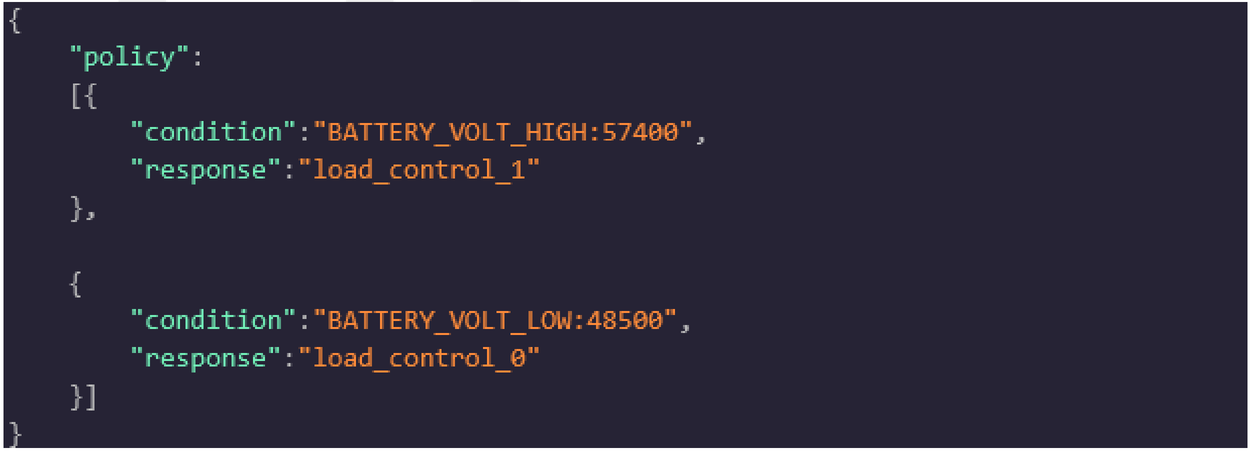

```
{
    "policy":
    [{
        "condition":"BATTERY_VOLT_HIGH:57400",
        "response":"load_control_1"
    },

    {
        "condition":"BATTERY_VOLT_LOW:48500",
        "response":"load_control_0"
    }]
}
```

Figure A-1. Example DCS Worker Node Policy

As an example of one worker node policy, depicted in Figure A-1 is a JSON formatted policy object that contains a list of "condition" and "response" pairs. The worker node parses this policy as: "When this condition is observed, enact this response." For this example, the two conditions are a set point for observed battery voltage. The responses are noncritical load control on and off. In this policy there are two situations described: (1) When battery voltage is above 57,400 mV, turn the load on; (2) when battery voltage is below 48,500 mV, turn the load off.

A.2 DCS Libraries

Table A-4. DCS Software Dependencies

Library	Language	Release	Module	Reference
go-libp2p	Golang	v0.30.0	Hive/Libp2p	https://github.com/libp2p/go-libp2p
libiec61850	C	v1.5.0	Worker/IED Communication Thread	https://github.com/mz-automation/libiec61850
libmodbus	C	v3.1.10	Worker/IED Communication Thread	https://github.com/stephane/libmodbus