CAISO Telemetry Solution Over Broadband Lab Test and Proof of Concept

Prepared by
Robert W. Anderson
Olivine, Inc.
&
Sam Piell
Pacific Gas and Electric Company

Date
May 2017

Acknowledgements

Pacific Gas and Electric Company's Demand Response Emerging Technologies Program is responsible for this project. It was developed as part of Pacific Gas and Electric Company's Demand Response Emerging Technology Program 2015-16 Portfolio. Olivine Inc. prepared this report with overall guidance and management from the Demand Response New Product Team at PG&E. Data analysis was conducted by PG&E's Data Analytics and Governance team. Lab testing was completed in PG&E's Emerging Grid Technologies lab. Both Olivine and PG&E would like to thank the staff of Rainforest Automation and Universal Devices for their collaboration on this study. For more information on this project, contact S4P5@pge.com.

Legal Notice

This report was prepared for Pacific Gas and Electric Company for use by its employees and agents. Neither Pacific Gas and Electric Company nor any of its employees and agents:

(1) makes any written or oral warranty, expressed or implied, including, but not limited to those concerning merchantability or fitness for a particular purpose;

(2) assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, process, method, or policy contained herein; or

(3) represents that its use would not infringe any privately owned rights, including, but not limited to, patents, trademarks, or copyrights.

Contents

1 Introduction 1
1.1 Background on CAISO Telemetry Requirements 1
1.1.1 When Telemetry is Required 1
1.1.2 Data Frequency and Quality Requirements 2
1.2 Project Overview 3
1.3 Architecture 4
1.3.1 PG&E SmartMeters™ 4
1.3.2 HAN Devices 5
1.3.3 Lab Scripts 6
1.3.4 Olivine DER 6
2 Lab Test Findings 8
2.1 Test Plan 8
2.1.1 Phase 1: Connectivity Testing 8
2.1.2 Data Collection 10
2.2 Measurement 11
2.3 Divergent Timing 12
2.4 Utilizing Summation Readings 14
2.5 Analysis of Accuracy 16
2.5.1 Per-location Accuracy 16
2.5.2 Missing Data 16
2.5.3 Accuracy 17
2.6 Project-Specific Challenges 19
2.6.1 PG&E's Emerging Grid Technologies Network Access Rules 19
2.6.2 Rainforest Eagle Challenges 19
2.6.3 Universal Devices ISY-994 Challenges 20
3 Findings 21
3.1.1 The CAISO 4-second requirement is not well supported by the configurations tested 21
3.1.2 Both HAN gateways are capable of providing 1-minute and 5-minute CAISO telemetry 21
3.1.3 The 1-minute instantaneous option has some benefits over the 5-minute option 21
3.1.4 Off the shelf devices require some configuration to securely connect to cloud systems 21
3.1.5 Missing and divergent data occurred with both devices in the tested solution 21
3.1.6 Potential Improvements to CAISO Direct Telemetry Rules 22
3.2 Future Research Questions 22
3.3 Conclusions 23

1 Introduction

PG&E's Lab Test and Proof of Concept of a *CAISO Telemetry Solution Over Broadband* was a 2016 project run by PG&E with Olivine and two hardware vendors to test the ability of existing Home Area Network (HAN) technologies to fulfill CAISO telemetry requirements. This report describes the following aspects of the lab study:

- Description of test plan
- Description of methodologies used
- Detail of lessons learned / challenges and barriers met within the Study
- Determine which, if any, of the methodologies support the CAISO BPM for 1 and 5-minute telemetry for PDR

The intention of this report is to summarize results of the project utilizing basic analysis and Olivine experience in wholesale market telemetry requirements, with the main purpose to inform whether the technology performs sufficiently well to merit additional study beyond the initial proof of concept, via a field study. The analyses described in this report are the best effort of the team to apply the CAISO Telemetry Requirements to the data collected. This report is meant to inform key stakeholders, including the CAISO and Demand Response Providers (DRPs) of one pathway under exploration for meeting the current telemetry requirements and increasing the ability to integrate Proxy Demand Resource (PDR) into the CAISO market.

1.1 Background on CAISO Telemetry Requirements

Telemetry requirements serve to ensure that CAISO operators have sufficient visibility to balance the real-time supply and demand on the transmission grid. The actual requirements vary by resource type, characteristics, and service provided. The C*AISO Business Practice Manual for Direct Telemetry* (BPM) describes the complete requirements for Direct Telemetry, including the responsibilities of various parties, and their roles as it relates to telemetry installation, validation and maintenance. For the purpose of this report, we focus on the data requirements, as that is what was tested.

1.1.1 When Telemetry is Required

Table 1 summarizes eligibility requirements for PDR participation in ISO wholesale markets. Most noteworthy, there is a requirement that any resource 10 MW or greater in size providing Energy, or of any size providing Ancillary Services must provide telemetry. The minimum load curtailment for participating in wholesale markets as a PDR is 0.1 MW for Energy resources and 0.5 MW for Ancillary Services resources.

PDR Type	Minimum Size for Which Telemetry is Required	Remote Intelligent Gateway	
		Maximum size for aggregation	Maximum size for single location
Energy	≥10 MW	≤ 400 MW & ≤ 25 Resource IDs	≤ 1200 MW
Ancillary Services (spinning/non-spinning reserves)	Always required		

Table 1: Eligibility Requirements for PDR Participation in Wholesale Markets

1.1.2 Data Frequency and Quality Requirements

Table 2 summarizes the CAISO data collection, transmission, and quality requirements for PDRs. Intervals for PDRs were originally closer to those that generators, but have been relaxed considerably for Energy and Non-Spinning Reserve.

Resource	Data Measurement Interval	Frequency at which RIG is queried for new data	Data Quality
Energy	5 minutes	every 4 seconds	+/- 2% of the true value[1]
Non-Spinning Reserve	1 minute		
Spinning Reserve	4 seconds		

Table 2: CAISO Measurement Interval, Frequency, and Data Quality Requirements for PDRs

The focus of the lab test was on the data measurement interval, and studying how those values would impact the data quality of an aggregated resource. The data measurement interval is described pictorially in Figure 1 for the case of 5-minute Telemetry, and characterized throughout the report as "polling frequency" as this is the term used to describe how frequently the HAN device is sending data to the RIG.[2]

[1] "True value" is defined as the accuracy of what is seen by the CAISO as compared to the value that a precision instrument would read consumption.
[2] Additional discussion of data measurement intervals can be found directly in the BPM or Appendix C of this report.

5-Minute Telemetry

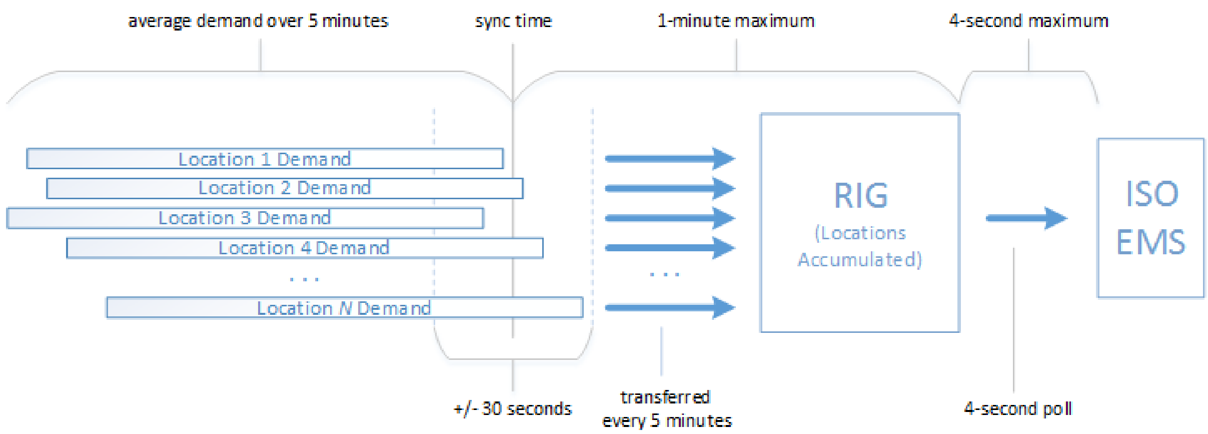

Figure 1: 5-Minute Telemetry for PDR[3]

The methodologies for calculating average demand over the 5 minutes are discussed in this report. At the heart of the requirements, whether 5 minute (Energy) or 1-minute (non-spinning AS), is ensuring that telemetry data reflects actual resource operations, in as near to real-time as possible.

1.2 Project Overview

PG&E's Demand Response Emerging Technologies (DRET) program undertook a lab study to test a telemetry solution set in a lab environment. This particular solution set was identified as having the potential to meet the business and technical objectives of CAISO telemetry, based on a survey of the technology landscape conducted in prior work funded by DRET[4]. It is viewed as DR provider-agnostic, meaning that it may be cost effective to implement, so that any eligible entity fulfilling the role of a DRP or providing services to such parties, would not have an institutional or commercial advantage over any other in providing its resource telemetry services. This is relevant, insofar as the data pathways are not run directly through the PG&E mesh network and can be leveraged by any DRP.

This effort demonstrated the viability of a particular architecture to meet CAISO's telemetry requirements so that Demand Response loads can participate in wholesale Energy and Ancillary Services markets. The architecture comprises these notable features, the combination of which distinguishes it from other options:

- A gateway appliance, paired to the SmartMeter™ via the ZigBee standard, and supporting the Smart Energy Profile (SEP) protocol for transporting demand and consumption measurements; the appliance interprets each incoming ZigBee data frame with an SEP payload, and forwards that in a device-specific formatted payload (based on standards or otherwise). The gateway provides a broadband connection to a cloud service, to be made available to a Remote Intelligent Gateway (described below), through a second interface, for exporting the telemetry data;

---

[3] From the CAISO Direct Telemetry BPM version 9. See https://www.caiso.com/rules/Pages/BusinessPracticeManuals

[4] Veregy Consulting for PG&E. *Assessment of Technologies Available to Meet California Independent System Operator (CAISO) Telemetry Requirements for PDR -- Final Report*. PG&E's Emerging Technologies Program, Project Number: DRET15PGE01. May 2016.

- A Remote Intelligent Gateway (or "RIG"), as defined by CAISO, for aggregating demand and consumption measurements arriving as telemetry from multiple SmartMeters™, and presenting the sums with the appropriate adjustments. The RIG is in communication with the gateway appliance over broadband.

The objectives of the project were to:

- Develop a testing methodology for HAN gateways to provide CAISO telemetry
- Test at least one such HAN gateway
- Determine suitability for such devices to support the various CAISO requirements for various products relevant to DR resources (i.e., energy, non-spinning reserves, and spinning reserves).
- To identify further research questions and efficacy of a follow on field-study.

1.3 Architecture

The overall architecture of the lab study is shown in the following figure:

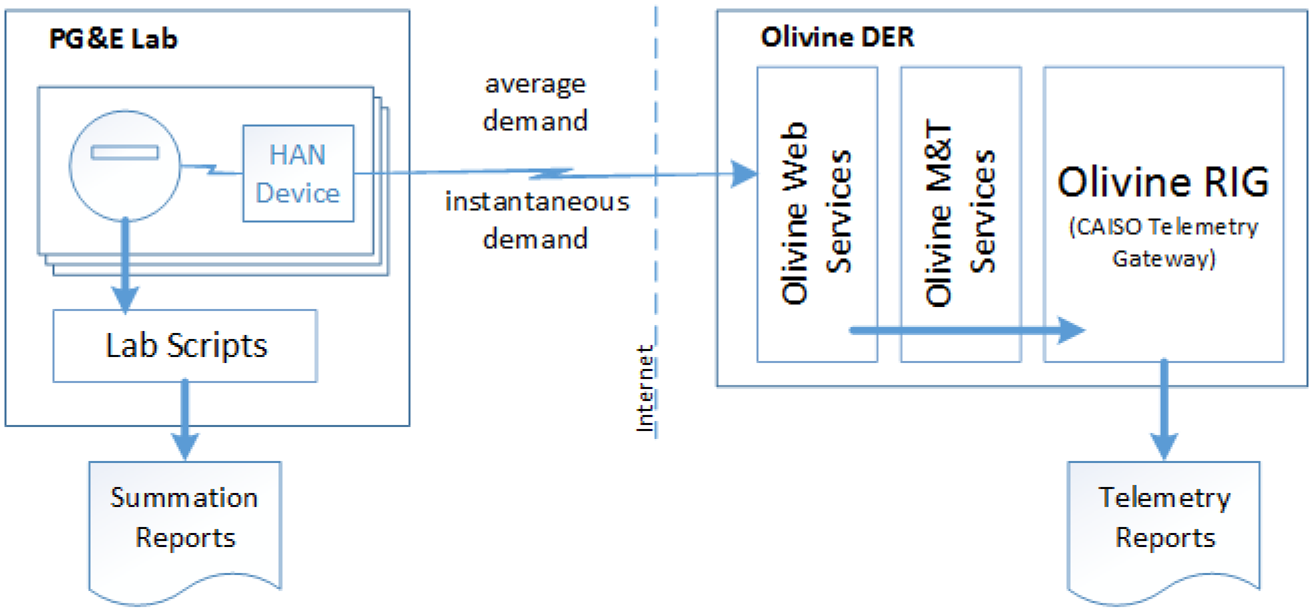

Figure 2: Study Architecture

This architecture is one in which the HAN device collects data from the SmartMeter™ and sends it directly to the Olivine DER system, the elements of which are described in the following sub-sections. Note that an alternate architecture was also contemplated, in which a HAN device manufacturer would collect data from their devices into their proprietary cloud with an integration between that cloud and the Olivine DER. This alternate architecture was rejected due to concerns of additional latency introduced by manufacturer cloud systems and the extra hop of data to Olivine DER.

The various components of this architecture are identified in the following sections.

1.3.1 PG&E SmartMeters™

A total of six utility-grade AMI meters reflecting a cross section of actual SmartMeters™ in production at PG&E were used in the project. These were the telemetry data source for measuring the instantaneous demand (i.e., kW) and energy consumption (i.e., kWh). The meters in use were three of each of the following:

- GE I210+ DSH @ 120 V

- L+G Focus DSH @ 240 V

In production, SmartMeters™ measure the power and usage for the customer's whole-premises. For the lab test, the measured load was generated from "load banks". These banks provided varying loads to the SmartMeters™ to establish accuracy of various telemetry frequency settings under different loads. The methodology is described further in Section 2.1.2.

1.3.2 HAN Devices

Devices from two manufacturers were included in the test. The devices were selected from the list of PG&E Validated HAN Devices, which meant that at some point, the device manufacturer had invested in a testing process to ensure compatibility with PG&E SmartMeters™.[5] [6]

Each HAN device was joined to a single PG&E SmartMeter™ and then communicated with that meter over the Zigbee protocol. In general, the HAN device would poll the meter for energy demand and usage and then push that data to the Olivine DER as shown in Figure 2. An important aspect of the HAN devices is that they do not measure demand nor energy usage. They rely on the SmartMeter™ for that; instead, they act as a gateway between the meter and the Olivine DER.

For both HAN devices, Olivine was responsible for altering the telemetry frequency with the each device having a different mechanism for performing that configuration change, identified in the following sections.

1.3.2.1 Rainforest Automation Eagle

The lab test included 6 Rainforest Eagle (Eagle) devices. This original model Eagle is a small form factor gateway designed to connect the HAN to the Internet. From the Rainforest Automation literature:

> The EAGLE™ Energy Access Gateway Link to Ethernet product is an Ethernet device that communicates directly with smart meters that have been equipped with ZigBee Smart Energy wireless capability. It functions as a gateway, providing smart meter data to the home Local Area Network (LAN) and the internet cloud.

The Eagles communicated with the PG&E SmartMeters™ using the SEP 1.1 profile, translating that data into an XML-based SEP-like formatted payload for sending to the cloud. While based on the SEP construct, this format is the Rainforest Automation-proprietary Uploader format.

This original model Eagle required firmware enhancements to support mutual authentication with the Olivine DER system utilizing client certification authentication as defined by the HTTPS protocol. Some issues that arose with the use of this device are covered in Section 2.6.2.

As a part of the lab test, Olivine procured client certificates for the Eagle devices to uniquely identify and authenticate each device against Olivine's DER. These certificates were issued by Comodo and were acquired by Rainforest using Web browser enrollment capabilities of Internet Explorer. Note that this mechanism allows the requester – Rainforest in this case – to create a private key and Certificate Service Request (CSR) while

---

[5] PG&E's Stream My Data program enables residential and small commercial customers to gain access to their own meter data in real time with the use of HAN devices. The HAN validation pre-existed this lab study in support of that program.

[6] Note that another device, Embertek's EmberPulse was considered as an additional device to test in the lab; however, at that time the lab test was focused on the "single hop" data flow (i.e., direct to Olivine, rather than via the manufacturer cloud) and Embertek was not able to support that.

requesting a corresponding public certificate from the issuing authority. This method of issuing certificates where the requester is the only party that holds the keys is considered a best practice[7].

The lab study required Olivine to configure changes to each device's polling frequency. The Eagle push frequency was modified by replying to pushed data with a Rainforest-proprietary response payload.

1.3.2.2 Universal Devices ISY-994

The lab test included 6 Universal Devices (UD) ISY-994 (ISY) devices. Another small form factor gateway, this device has a much broader feature set. From the UD literature:

> ISY994 Series energy management and automation controller provide power and flexibility perfect for middle market, small, and medium business needs. Using open standards, the ISY994 provides building owners, contractors, and project leads the flexibility to automate a variety of every day tasks. [It is] capable of communicating with a wide variety of off-the-shelf devices such as thermostats, lighting, IoT devices, and loads using ZigBee, Z-Wave, or Insteon, in addition to IP and direct contacts. The ISY994 series can read energy consumption from individual devices using Z-Wave, ZigBee, and AMI smart meters and report the data using OpenADR . . .

The ISY devices were "off the shelf" with no project-related firmware updates and utilized the OpenADR 2.0b protocol. The devices did still need to be correctly configured with OpenADR Alliance VEN (Client) certificates, with the Olivine VTN server information, and some specific OpenADR settings to enable telemetry push through the OpenADR EiReport endpoint.

The ISY-994 push frequency was defined through the OpenADR EiReport service methodology. Olivine utilized the features of OpenADR 2.0b to configure such reports to facilitate the project.

1.3.3 Lab Scripts

PG&E's Emerging Grid Technologies lab was responsible for implementing a software program or "script" that retrieved data directly from the six PG&E SmartMeters™, via a specially established local test environment on the Silver Spring Network platform. The purpose of the scripted program was to retrieve usage information from the meters for the purposes of comparing and establishing the accuracy of the data provided by the HAN devices to Olivine. The output of these scripts was provided to PG&E for analysis. The use of these scripts is described further in Section 2.5.1.

1.3.4 Olivine DER

The Olivine DER is an OpenADR 2.0b certified flexible cloud software system designed to manage the operations around retail and wholesale demand response and distributed energy resource programs. Several specific features of Olivine DER are identified in Figure 2:

Olivine Web Services:	Provides API access to various features of Olivine DER, including the Rainforest Automation Uploader API and OpenADR HTTP access.
Olivine M&T Services:	Provides various metering and telemetry services, including aggregation of such data for the purposes of operations and settlements.
Olivine RIG:	The CAISO-validated Remote Intelligent Gateway (RIG) that is the basis for market resources – including DR resources – to provide real-time telemetry to the CAISO.

[7] Retaining control over private keys is a critical component of key management. See NIST Special Publication 800-57 *Recommendation for Key Management*, Chapter 5.

The Olivine DER offers multiple pathways to receive location-level telemetry, noting that two unique methods were used in the lab study. For the Rainforest Eagle, Olivine added the ability to receive data from the Rainforest Automation Uploader API. In this case, data were pushed to Olivine DER using a simple HTTP push API with XML payloads. In the case of the UD ISY-994, all telemetry was provided to Olivine DER using the OpenADR 2.0b EiReport endpoint, again using XML payloads. In both cases, mutual authentication was performed utilizing client and server certificates.

Under normal conditions of providing telemetry for a market resource, the CAISO EMS would utilize DNP3 over a secure channel to receive telemetry points from the Olivine RIG. As an alternative to this approach, and to meet the needs of the lab study, Olivine provided Telemetry logs to PG&E to perform the analysis.

2 Lab Test Methodology and Data

2.1 Test Plan

The lab tests were split into two main phases. The first phase was focused on connectivity of the devices themselves. This included HAN standards compliance, proper authentication, and encryption to ensure overall end to end behavior to meet the needs of the project. The second phase of the tests was focused on telemetry data collection for analysis purposes.

2.1.1 Phase 1: Connectivity Testing

This phase of the study was focused on end-to-end connectivity of all of the components as shown in Figure 2: Study Architecture. This included the following tests, described in the following sections:

- ZigBee-SEP Compliance
- Authentication and Encryption
- Gateway Power Cycling and Broadband Service Interrupt
- Frequency Configuration
- End-to-end Behavior

2.1.1.1 ZigBee-SEP Compliance

The Rainforest Automation Eagle had updated firmware to comply with the project requirements, and as such was checked for ZigBee-SEP compliance using a relevant subset of tests focusing on the basic connectivity features of the protocol. These tests were run at the PG&E lab and included the following tests:

- Time Synchronization with the Meter ESI, as defined in ZigBee 07-5356-18
- Certificated base key establishment with the Meter ESI as described in ZigBee 07-5356-18
- Secure rejoin on power outage
- Rejoin to a different Meter ESI, and establish a key based on the certificates

The Universal Devices ISY-994 had already received a compliance certification from PG&E and therefore additional compliance testing was deemed unnecessary.

2.1.1.2 Authentication and Encryption

The purpose of this test was to ensure that the devices were securely authenticated to the Olivine DER and that data was sent in a confidential manner. The test itself was simply to configure the devices with proper credentials and the Olivine DER endpoint and then allow the devices to push data to that endpoint.

A successful result was the confirmation that Olivine DER's ability to log device payloads from the device. Note that this result confirms both authentication and encryption since the transport-level security (i.e., HTTPS relying on TLS) in force by Olivine DER requires encryption and mutual authentication. Non-encrypted data could not be logged because it wouldn't be accepted by the Olivine server, though in practice a client negotiating TLS with a server would never transport the payload data (encrypted or otherwise) without a successfully negotiated encrypted channel. This was the case both for the Rainforest Uploader API and the OpenADR 2.0b protocol used by the UD ISY-994.

Note that issues did arise with the Rainforest Eagle regarding authentication – covered in Section 2.6.2 – however, as expected, unsecured information was not passed to nor accepted by the Olivine DER in this circumstance.

2.1.1.3 Gateway Power Cycling and Broadband Service Interrupt

These tests were intended to ensure that the gateway devices were able to continue normal operation after a loss of power or broadband service interrupt.

To test the ability of the gateways to correct themselves after a loss of power, the devices were unplugged and then plugged in again to see if they would proceed to submit data. To test the ability of the gateways to correct themselves after a broadband outage, the devices' Ethernet cords were disconnected and then reconnected to see if they would proceed to submit data once back online.

The power on/off and disconnect/reconnect times were reported by the PG&E lab, while the data received times were identified by Olivine (with the times corrected to match the 1-2 seconds difference between the PG&E lab and Olivine clocks noted at the time).

	Rainforest Eagle (4/22/2016)	UD ISY-994 (8/26/2016)
Powered Off	10:14:50	09:10:00
Powered On	10:15:00	09:11:00
Data Received	10:36:05*	09:11:46
Bandwidth Disconnected	10:08:00	10:05:00
Bandwidth Reconnected	10:10:35	10:06:00
Data Received	10:10:35	10:06:31

Table 3: Power and Connectivity Tests

*Note that after power on, the Eagle took over 20 minutes to begin sending data to Olivine DER. See Section 2.6.2 for more information and the resolution to that issue.

2.1.1.4 Frequency Configuration

This test was to prove both the ability for Olivine to remotely configure the devices to alter the telemetry push frequency and that the devices would be responsive to this change.

The two devices have different configuration options, so the specifics for changing configuration were different (see Section 1.3.2); however, the test followed the same approach:

1. Turn on logging on the Olivine DER to set a baseline for the frequency of pushed data.
2. Modify the configuration for frequency
3. Examine the logs to determine if the change in configuration was accomplished.

Ultimately these tests were deemed successful as proven out in all following tests.

2.1.1.5 End-to-end Behavior

This last portion of connectivity testing was intended to determine a lower bound for communication from the devices. Early discussions with PG&E lab personnel made it clear that the PG&E SmartMeters™ are not capable of providing HAN data more frequently than approximately every 10 seconds. To confirm this, frequencies between 4 seconds and 5 minutes were tested.

The Rainforest Eagle was configurable to a 4-second interval. It was fairly consistent – barring some missing intervals – of pushing data every 4-seconds; however, interestingly the instantaneous kW that it provided did not change more often than once every 4 payloads. That is, the kW reports would be the same for at least any 16-second period even when the kW values changed more frequently. While this was not a lengthy analysis, this may indicate a limitation of the meter's ability to provide instantaneous data to meet the 4-second requirement. Alternatively, it could imply that the Eagle is not transmitting data correctly; however, it is Olivine's understanding, that the Eagle only transmits data when it receives that data from the meter in response to a SEP 1.1 poll. This would imply that the meter is providing the same value across time. See Appendix A for an example.

During subsequent tests using the ISY-994, the highest frequency tested was at 10-seconds. At this frequency, the ISY-994 was able to consistently deliver payloads. In the case of the ISY, changes to instantaneous readings were never more frequent than every 30 seconds, even when the kW values changed more frequently. This may have been a configuration issue, and merits additional testing.

2.1.2 Data Collection

This phase of the lab study tested the HAN devices utilizing varying loads on the respective SmartMeters™ over multiple hour periods.

Two profiles for varying loads were selected, as shown in Figure 3. These profiles were programmed into load banks alternating between each profile every hour, aligned with the "top of the hour". So, for example, if the load banks were started at 12 PM, the first hour would be Profile A, and the next hour would be Profile B, then back to Profile A, etc., until the load banks were turned off. This approach was selected to decouple the activities of PG&E Lab staff from the reconfiguration necessary to change telemetry frequency by Olivine, which allowed reduced orchestration between these two parties, thus shortening the timeline to complete the tests.

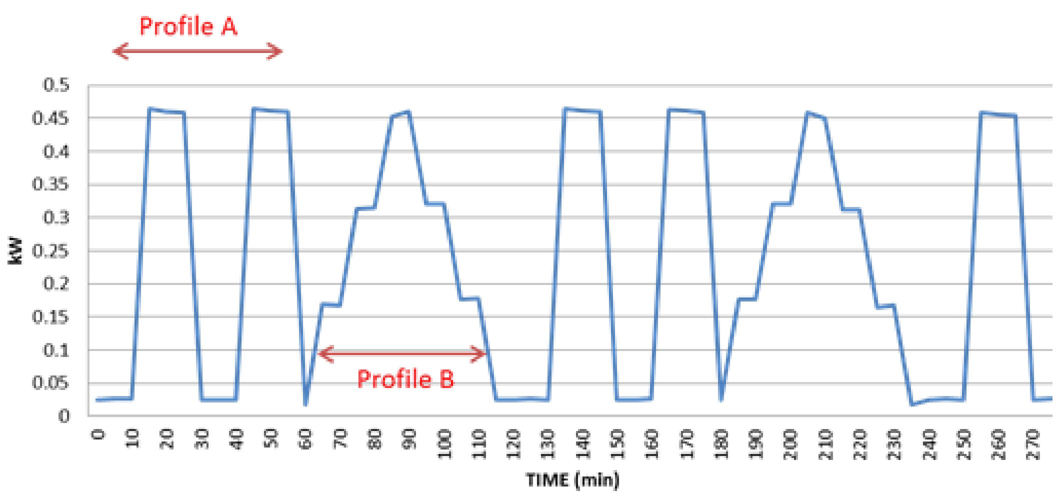

Figure 3: Load Bank Profiles

The primary purpose of varying loads was to determine if different levels of load had any impact on the frequency of update to the summation reading from the meter. The testing did prove this out, with lower load values taking longer to register as changes in the summation readings. See the tables in Section 2.4.

The data collection was performed over several different periods. This was to shift between the two manufacturer's devices – due to a limitation in the number of devices that could be installed at a time – and was also due to incremental learning from the initial data collection efforts.

The initial data collection and analyses on the lab tests were focused on telemetry accuracy and latency from the perspective of understanding the limitations of the devices to provide high-frequency data. Subsequent data collection and analysis focused optimizing the actual HAN frequency to best meet the CAISO requirements.

Testing of the Eagle for 10, 20, 30, 60, 150, and 300 seconds was performed on July 20, July 21, July 27, August 1, August 9, August 10, and October 21. In addition, tests were performed continuously from October 24 through October 31 running at 300 seconds.

Testing of the ISY-994 was accomplished in two steps. In a single 24 hour period from August 18 - August 19, tests were performed at 10, 20, 30, 60, 150, and 300 seconds. The second step was from October 20 through October 23 at 300 seconds.

Note that there was a larger amount of testing performed on the Eagle than the ISY-994. This was partially due to ramping up the lab test and standardizing the reporting, and also due to retest required while working through issues with the Eagle.

The complete list of time periods is available in Appendix B.

2.2 Measurement

PG&E SmartMeters™ measure instantaneous power (in watts) and accumulate total energy used over time in watt-hours. These meters also record intervals of energy usage. In the residential metering context, the recording interval tends to be 1 hour. In addition to these hourly usage intervals, the meters persistently store an indicator of total energy usage (i.e., the summation) measured over the life of the meter. Utilizing the Zigbee radio interface to the meter, only two of these electrical measurements are readily available and applicable to telemetry: the instantaneous power and the total energy usage. The other element – the recording interval reads – is not relevant because the interval being recorded will generally be too long to be useful for telemetry[8].

There are two main approaches to meeting the CAISO telemetry requirements. The first is to utilize instantaneous demand. This measurement is straightforward because it comes directly from the meter through the HAN without interpretation. Because it is merely a snapshot of the demand at a single moment, this is best suited for the two higher-frequency options for telemetry (i.e., for 4-seconds or 1-minute). The second approach is to utilize average demand over a 1-minute or 5-minute interval[9]. Noting again that the meter itself cannot provide this calculation, another component of the telemetry system architecture needs to compute this average demand value.

To compute average demand, there are two methods that could be utilized given the data available from the HAN: the first is to make high-frequency instantaneous reads, perhaps at the sub-second level, and average them over an interval of time. In other words, for a 5 minute interval, the instantaneous (kW) readings could be taken every one minute and averaged together. One shortcoming with this approach is that for intervals of one minute or less, it would rely on consistent higher-frequency data at levels faster than can be delivered through

[8] For PG&E metering there is an intersection at 5-minute intervals between the highest frequency metering by PG&E and the lowest-frequency telemetry allowed by CAISO for energy-only PDRs; however, 5-minute metering at PG&E is generally performed at large customers where the use of HAN gateways has not yet been established.
[9] Note that the CAISO Direct Telemetry Business Practice Manual does not dictate when instantaneous demand or average demand reading should be used.

the Zigbee interface, which as identified in Section 2.1.1.5, cannot be relied on. The second method is to utilize the summation reads and determine the energy used between reads, computing an average as follows:

$$inferred\ average\ load = \frac{\Delta summation\ (kWh)}{\Delta time\ (hours)}\ kW$$

where:

$$\Delta summation = summation_i - summation_{i-1}$$
$$\Delta time = time_i - time_{i-1}$$
$$i = index\ of\ a\ reading$$
$$i-1 = index\ of\ the\ previous\ reading$$
$$summation_i = total\ usage\ for\ reading\ i$$
$$time_i = time\ that\ reading\ i\ occurred$$

If summations are continuously accurate, and accurately associated with a time, then this is a theoretically perfect system to infer average demand; however, that is not the case, so there are limitations to this approach:

- Discussions with various team members and direct experience retrieving summation from the meters shows that summations do not continuously update. In other words, the kWh readings of the meter update with discrete values rather than continuously
- The time of the summation reading is the time that the Zigbee radio received the payload from the meter, not the time at which the meter updated the summation.
- Exact details on these two items are likely implementation-specific based on HAN devices and specific firmware revisions.

2.3 Divergent Timing

In the subsequent section, this report will focus on the accuracy of telemetry data for the purposes of meeting CAISO requirements; however, the focus of this section is on the reliability of receiving data from the devices on a set period. This is of interest in understanding the behavior of these devices and could be an input into future recommendations for updates to CAISO requirements. The team used the term *divergent data* to describe any readings that were outside of the expected frequency, both early and late. Two main criteria to determine divergence were utilized: the exact period expected, and using a T-test to define outliers at 95% confidence.

Table 4 shows the percentage of divergence based on these two criteria for different frequency settings. p is the expected period of data receipt; Δt is the actual time between two payloads received from a single device. Note that time units are in seconds.

		Divergent Percentage	
	Divergent Criteria	EAGLE	ISY
$p = 10$			
Exact period	$\Delta t \neq 10$	80.26%	7.18%
T-test 95%	$\Delta t < 7 \cup \Delta t > 13$	12.81%	2.44%
$p = 20$			
Exact period	$\Delta t \neq 20$	74.66%	14.76%
T-test 95%	$\Delta t < 15 \cup \Delta t > 25$	3.49%	8.64%
$p = 30$			
Exact period	$\Delta t \neq 30$	84.15%	8.01%
T-test 95%	$\Delta t < 24 \cup \Delta t > 37$	6.61%	6.34%
$p = 60$			
Exact period	$\Delta t \neq 60$	71.24%	15.81%
T-test 95%	$\Delta t < 54 \cup \Delta t > 68$	2.95%	8.68%
$p = 150$			
Exact period	$\Delta t \neq 150$	84.31%	11.28%
T-test 95%	$\Delta t < 143 \cup \Delta t > 200$	3.28%	3.35%
$p = 300$			
Exact period	$\Delta t \neq 300$	76.78%	8.76%
T-test 95%	$\Delta t < 260 \cup \Delta t > 340$	0.84%	2.68%

Table 4: Divergence Rates

Although the team hypothesized there would be a consistent trend of decreased divergence as *p* increased, this was only partially true. While both devices trended toward decreased divergence when comparing the 300 second to the 10 second interval, it was not a consistent drop. In some cases, the divergence actually increased. Figure 4 illustrates the divergence rate trend over time.

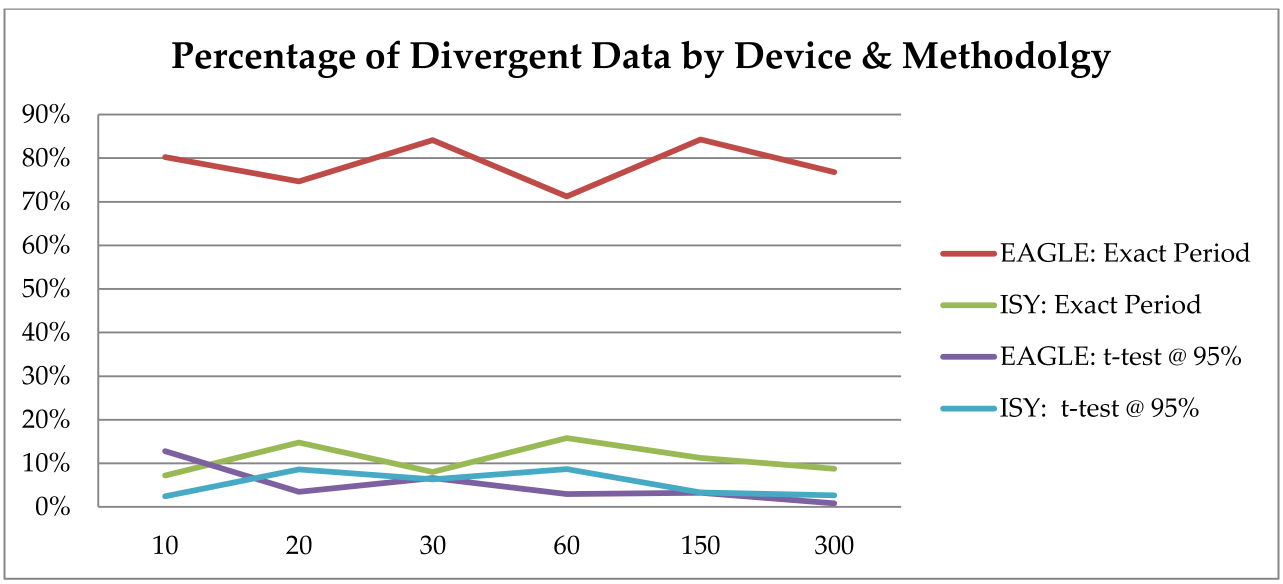

Figure 4: Illustration of Divergence Trend by Device and Precision Level

The following charts show the distribution of time difference between each reading for both the Eagle and ISY devices set for a 1-minute and 5-minute period. The red line indicates the expected period, which is 1-minute and 5-minutes. As clearly shown, on average about 50% of the data was sent earlier than the designed interval.

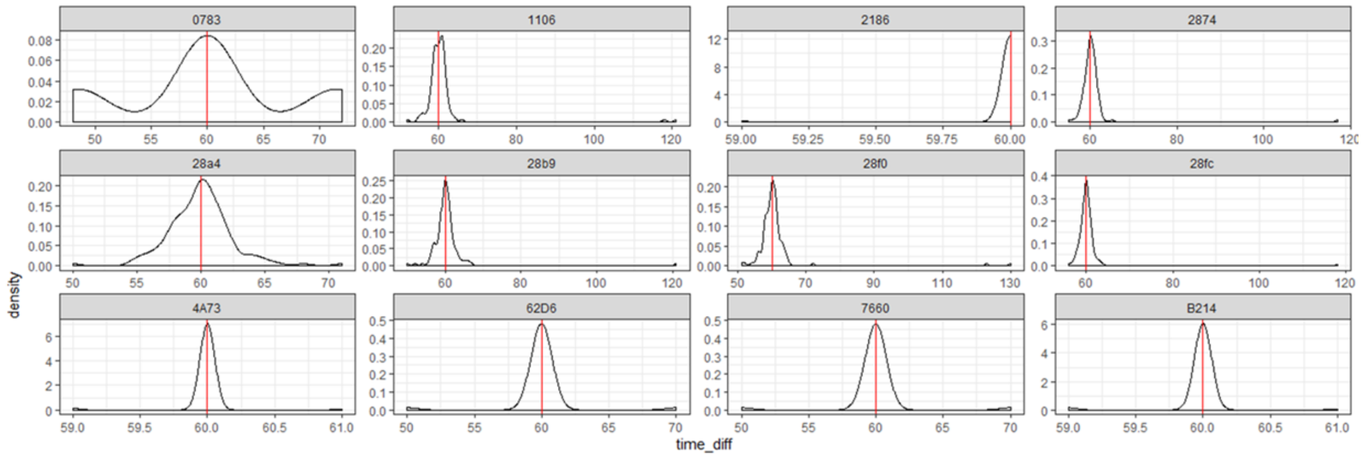

Figure 5: HAN device 1-minute timings

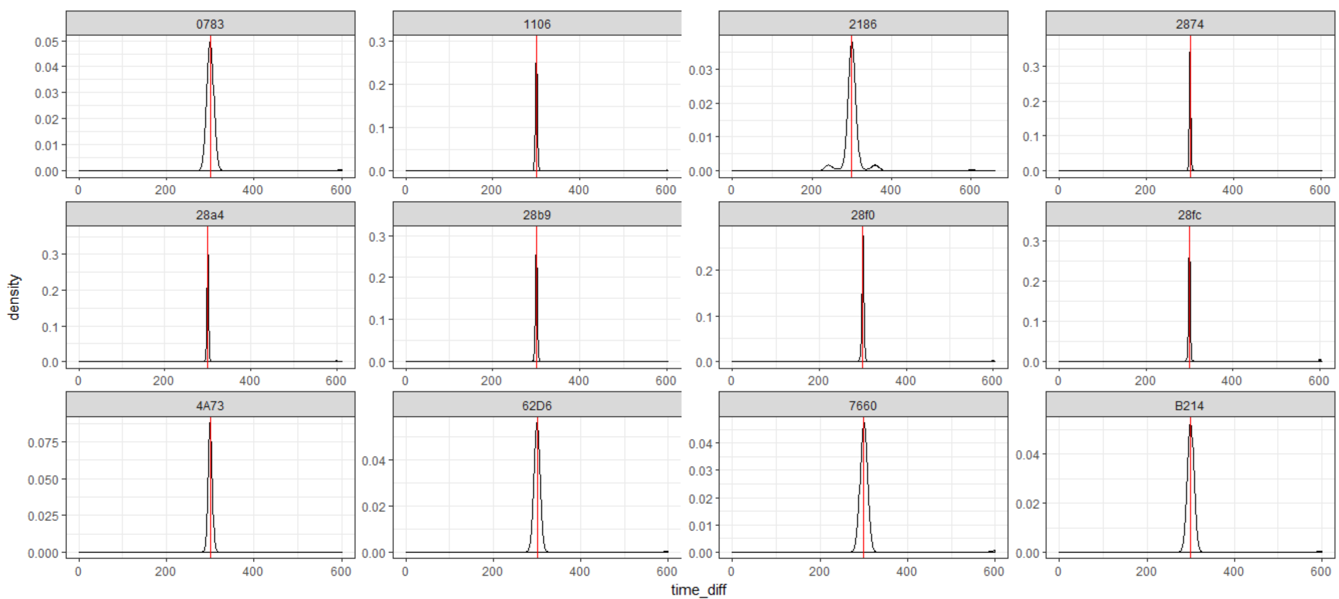

Figure 6: 5-minute timeliness

2.4 Utilizing Summation Readings

One of the questions this study attempted to answer was what is the impact of the load amount and frequency of change in that amount on meter-reported summation. The tests show that at very low loads (e.g., .033 kW), the summation value appears to change less frequently. This may come down to a rounding issue (i.e., the fidelity at which the meter reports summation); however, the end result is that at low levels of energy, using the inferred average load becomes less useful than using instantaneous readings.

The example in the following tables shows this very clearly, comparing one extreme .033 kW collected once per minute in Table 5, and .95 kW collected every 10-seconds in Table 6. This is from the same meter and same device at different times of the day. What we see is that the Summation value changes approximately once every two minutes at .033 kW; however, it changes approximately once every 20 seconds at .95 kW.[10]

Time	Load	Summation	Inferred Average Load
16:02:03	0.034	26.685	0.059427
16:03:16	0.033	26.685	0
16:04:02	0.033	26.686	0.077683
16:05:00	0.033	26.686	0
16:06:03	0.033	26.687	0.056679
16:07:03	0.033	26.687	0
16:08:01	0.033	26.688	0.061836
16:09:09	0.033	26.688	0
16:10:00	0.033	26.689	0.070523
16:11:02	0.033	26.690	0.057744
16:12:01	0.032	26.690	0
16:13:03	0.033	26.691	0.058006

Table 5: Summation changes at 1-minute Frequency

Time	Load	Summation	Inferred Average Load
6:43:22	0.95	9820.833	2.083174
6:43:33	0.95	9820.836	0.96268
6:43:43	0.95	9820.836	0
6:43:56	0.95	9820.84	1.107688
6:44:05	0.95	9820.84	0
6:44:12	0.95	9820.844	1.936135
6:44:22	0.95	9820.848	1.469837
6:44:31	0.949	9820.848	0
6:44:42	0.952	9820.852	1.407018
6:44:55	0.952	9820.856	1.115745
6:45:02	0.952	9820.856	0
6:45:12	0.95	9820.860	1.486286
6:45:21	0.95	9820.864	1.479469
6:45:31	0.952	9820.864	0

Table 6: Summation Changes at 10-second Frequency

Note that the “0” for “inferred average load” implies that the summation value did not change. Therefore, we conclude that this could have implications for the accuracy of measuring small load drops on a single meter.

[10]Although not shown, the “inferred average load” for a lower load at the 10-second frequency would have even more “0” results.

2.5 Analysis of Accuracy

The goal of the analysis is to understand how well the devices operate to assess if they can meet the requirements set forth in the CAISO BPM, described in Appendix C. Particularly key is the accuracy requirement for telemetry data. The California ISO BPM on Direct Telemetry states that “All telemetry data reported via the RIG must be within +/-2% of the true value”[11]. There is no further direction on how to apply this to an aggregation of locations. Much effort was put into this exact question in the Supply DR Integration Working Group completed in 2015, and while that work was adopted in principle by the CAISO, it was never codified in the BPM. See Appendix D for an excerpt of the working group report.

The core issue is that +/- 2% requirement may seem straightforward, but there are open questions as to how to define and meet this requirement in a potentially large, distributed aggregation of underlying locations. The standard itself was defined specifically to be more relaxed that the +/- 0.2% requirement for wholesale metering, to enable less expensive measurement with outboard current transformers versus self-contained utility-grade meters.

In the absence of direction from CAISO, this paper identifies an approach for establishing accuracy. The approach involves understanding the latency, accuracy and reliability of the devices and communication pathways, as data travels from the SmartMeter™ to the RIG. PG&E SmartMeter™ data was used as an anchor of comparison for part of the data analysis and the data actually received by the RIG compared to the device frequency setting was used to define missing data. Ultimately, this goal of +/-2% accuracy, inclusive of latency and reliability was the target and is described in Section 2.5.3. Note that the rest of this section focuses on the 1-minute and 5-minute data collection because those fit in with the CAISO requirements.

2.5.1 Per-location Accuracy

The accuracy of the underlying locational telemetry data comes down to two parts: the accuracy of the SmartMeter™ and the accuracy of the HAN device to report that data.

The PG&E SmartMeters™ are primarily Landis & Gyr or Aclara/GE meters which are either accurate to +/- .2% or +/- .5%. For the purposes of Section 2.5.3, the more conservative number is used.

The HAN gateways themselves should simply be passing meter measurements through to the Olivine DER and do not measure power or energy. Because it is conceivable that incorrect values could be reported or data missing altogether, PG&E implemented scripts in the lab that retrieved summation readings from the meter directly and logged that data for comparison against data provided to Olivine DER from the gateway devices.

Once this PG&E logging system was testing and running correctly, comparisons against the Olivine DER logs were 100% aligned, barring occasional discrepancies related to missed or dropped readings.

As a result, the per-location accuracy was deemed to be equivalent to the meter, with +/-.5% used as the most conservative number.

2.5.2 Missing Data

Considering that any concerns about accuracy of individual telemetry metering were resolved to the team’s satisfaction, the team turned to identifying the completeness of the data pushed to Olivine DER from the HAN devices. While Section 2.3 focused on divergence, the primary concern here is that of missing data. In this case, we focus on the 1 and 5-minute cases only since these are directly applicable to CAISO telemetry requirements.

[11] CAISO Business Practice Manual for Direct Telemetry version 10.0, Section 5.5 Data Validation and Confidentiality

The concept of *missing data* is when an expected payload does not arrive on or before the expected time, it is considered as missing. It was expected that such a situation would occur generally due to some latency; however, the project logs showed that such data could be significantly late such that it could not be explained by network latency (e.g., 2 minutes late for a 5-minute frequency). The more common occurrence of missing data was from an entirely dropped data packet or one that was just not sent.

The following table shows the percentage of missing readings based on the additional two criteria for different frequency settings. p is the expected period of data receipt; Δt is the actual time between two payloads received from a single device. Note that time units are in seconds.

		Missing Percentage	
	Missing Criteria	**EAGLE**	**ISY**
$\boldsymbol{p = 60}$			
Exact period	$\Delta t > 60$	36.89%	7.13%
T-test 95%	$\Delta t > 68$	1.26%	3.92%
$p + 60$	$\Delta t > 120$	0.56%	0.00%
$\boldsymbol{p = 300}$			
Exact period	$\Delta t > 300$	39.30%	3.68%
T-test 95%	$\Delta t > 340$	0.77%	1.73%
$p + 60$	$\Delta t > 360$	0.77%	0.91%

Table 7: Missing Rates

We can see a fairly high missing percentage if the period *p* is treated as a hard cut off for new data. Using a T-test to define outliers at 95% confidence, the percentage of missing data drops substantially if an additional 8 seconds are added to the one minute frequency and 40 seconds to the 5 minute frequency,. Finally, because of the 1-minute grace period provided in the CAISO BPM on delivery of data as covered in Section 2.5.3, the *p + 60* results provide relevant values to answer questions about accuracy as they would be applied in support of the CAISO telemetry requirements.

2.5.3 Accuracy

In the previous sections we have identified the accuracy of the per-location telemetry and the incidence of missing data from the HAN devices.

The implications of this accuracy for instantaneous readings are clear: the per-location telemetry reported is exactly that reported by the meter, and therefore +/- .5% accurate as described in Section 2.5.1. For readings that are inferred from the changes in summation, this becomes a little less clear; however, on average that methodology which averages the change in summation over adjacent time intervals, will result in the same accuracy as the underlying metering. Further investigation and analysis is warranted, however, before we recommend this method considering its reliance on certain behavior of the SmartMeter™. If we accept this premise, then both instantaneous and inferred readings provide +/- .5% underlying measurements.

Next we have to account for missing data as described in Section 2.5.2 and confront what this means for the entire aggregation. While not explicitly stated, it is a reasonable interpretation that the CAISO accuracy requirement applies to the sum of the all underlying locations. The challenge comes in resolving the issues when telemetry readings are missing for one or more underlying locations[12]. There are various ways to describe accuracy in this context.

For example, imagine an aggregation of 100 locations, with all telemetry devices reporting at +/- .5% accuracy. If 1 device stops reporting data, and that value is replaced with a zero, then one might argue that the accuracy has been reduced by 1%. That may be a reasonable definition of accuracy, but it is only fully correct in the case that the 100 locations are reporting uniform demand data. Otherwise, the impact may be different. For example, in a case that the 1 "missing" location has an average demand 99 times higher than every other single, uniform load. In that case, the accuracy has fallen to 50%. The reality lies between these two cases of perfect uniformity and a very disparate population.

A more realistic case might be where the loads are more predictable, and that missing data is estimated with a +/-5% accuracy. In the first example of uniform demand data when 1 device stops reporting data, we now have a much better outcome:

$$\left(\pm 5.00\% \times \frac{1}{100}\right) + \left(\pm 0.50\% \times \frac{99}{100}\right) = \pm 0.56\%$$

And, following this through, with uniform demand data, we can get up to 33% of missing data without crossing the 2% threshold:

$$\left(\pm 5.00\% \times \frac{33}{100}\right) + \left(\pm 0.50\% \times \frac{67}{100}\right) = \pm 1.99\%$$

No real world scenario would have such uniformity of data and further data collection and analysis may be warranted to investigate this further, noting that these examples suggest that estimation techniques and some uniformity of customer type and usage patterns may be required for aggregated resources to meet accuracy requirements.

Bringing this back to the lab study, we achieved the following, assuming a +/-5% accuracy of the estimation when there are missing data:

	Eagle	
Period	Missing %	Accuracy %
1-minute	0.56%	$(\pm 5.00\% \times 0.56\%) + (\pm 0.5\% \times (1 - 0.26\%) = \pm 0.53\%$
5-minute	0.77%	$(\pm 5.00\% \times 0.77\%) + (\pm 0.5\% \times (1 - 0.81\%) = \pm 0.54\%$

Table 8: Eagle Accuracy

	ISY	
Period	Missing %	Accuracy %
1-minute	0.00%	$(\pm 5.00\% \times 0.00\%) + (\pm 0.05\% \times (1 - 0.26\%) = \pm 0.50\%$
5-minute	0.91%	$(\pm 5.00\% \times 0.91\%) + (\pm 0.05\% \times (1 - 0.81\%) = \pm 0.54\%$

Table 9: ISY-994 Accuracy

As shown in these tables and based on these assumptions, the devices met the CAISO accuracy requirements.

[12] Outages in a distributed system such as this, but particularly when leveraging a customer-installed HAN device that relies on other home networking components and broadband connectivity.

2.6 Project-Specific Challenges

There were several project-specific challenges that impacted the schedule of completion, but were not relevant to the central question of suitability for such HAN devices to provide CAISO telemetry. As such, they are identified here but are not considered "lessons" of the project. This includes configuration of the devices themselves, because at any scale all configuration would be performed in bulk by the manufacturer or project implementer or even over the air in some configurations.

2.6.1 PG&E's Emerging Grid Technologies Network Access Rules

As can be expected, the PG&E Emerging Grid Technologies lab has strict and defensible requirements for network security. In particular, this default security places limitations on access to devices under test within the lab. For example, Rainforest Automation has a method to perform remote diagnostics and over the air firmware updates. These capabilities require network firewall exceptions to be in place and approved by senior lab personnel. During the design of the project and discussions with the lab, a firewall exception was deemed too challenging to acquire for this project and so was not undertaken. As a result when challenges with the Eagle arose, there was no fast way for Rainforest to diagnose and fix the problem. The two solutions that were utilized were to remove devices from the lab environment – literally to bring them into the home of lab personnel – or to ship devices back to the corporate office in Vancouver, Canada. This added significant time to the project schedule, and may be viewed as a general trade-off of security versus ease.

While it is not a central theme of this report or of the lab study, it would be helpful if there were a simple mechanism to provide access to such devices in the lab. One solution that would have worked for this project would be an alternate secure VLAN with a path to the Internet with firewall protections similar to a home network configuration. This could be particularly useful and important as the lab increases the frequency of testing devices for the Internet of Things.

2.6.2 Rainforest Eagle Challenges

As identified in Section 1.3.2.1, Rainforest Automation enhanced the Eagle gateway to support authentication with client certificates. There were several challenges that occurred during the project related to the Eagle gateway. Aside from authentication issues that were clearly related to the firmware enhancements, Olivine cannot confirm if the other issues were also related.

During the connectivity-testing phase of the project, these challenges arose:

- Administrative user configuration to set endpoint URL to the Olivine DER was not being honored.
- Device restart could result in devices being offline for 45 or more minutes.
  - It was confirmed by Rainforest that when Uploader API requests were set to be infrequent, then the devices would demonstrate a significant amount of time between applying power to first transmission to Olivine DER. For example, when the summation request was 5-minutes, the startup time could be 45 minutes.
  - This issue was only apparent when the Eagle was set to synchronize time with the meter instead of an Internet time server and was pushing data to a remote server protected with HTTPS as Olivine DER. The dependency here was that Rainforest had reduced the frequency of time synchronization to improve meter data rates for the measurement data required for the project. The result was that until time was synchronized, the Eagle could not verify the validity of the remote server certificate, and therefore data "push" would fail. This issue was corrected by Rainforest during the project. Note that this was a custom configuration completed by PG&E personnel and not "out of box" behavior.
- Some of the polling performed between the device and meter and data sent to the Olivine server cannot be controlled by Olivine. This resulted in additional 'chattiness' on the HAN and unnecessary data

pushed to Olivine DER. Some of the additional data includes HAN connectivity status, pricing data, and time events, none of which were actionable for this project. In addition, since this data is also exchanged with the SmartMeter™ there was some evidence that this additional data impacted the reliability of the device at high-frequencies. While this data does have applications, for the purposes of delivering CAISO telemetry it could be helpful to disable it completely. In fact, doing so might allow for an increase in reliable frequency, though it is unclear if reliable 4-second telemetry could be achieved due to meter limitations.

During the data collection phase, these challenges arose:

- While in Connectivity Testing the authentication enhancement appeared to work correctly; however, in the data collection phase issues became evident where a high percentage of payloads were being rejected by the Olivine DER with an HTTP 401 Unauthorized Status. This error was evident in the Olivine logs as well as, on inspection, logs on the Eagle devices. This was a case where the Eagles had to be shipped back to Rainforest Automation for diagnosis and ultimately to be fixed.
- A special case of missing data was uncovered in that the Eagle pushed "empty" data packets to the Olivine DER. The actual condition was a receipt of the appropriate XML but without any values within the XML tags. For example, a normal payload might contain `<Demand>0x000020</Demand>` to indicate instantaneous demand, but the empty payloads contained `<Demand></Demand>`.
    - This condition was evident in the Olivine DER logs identified as rejected payloads due to invalid content.
    - Rainforest stated that this case could only occur if the SmartMeter™ were to provide incomplete or invalid data back to the Eagle. As a result, the PG&E lab captured Zigbee traffic to test this assertion. These logs combined with Olivine DER Web service logs from the same period showed that this hypothesis was incorrect: that is, the SmartMeter™ was providing valid responses, but the Eagle was not providing valid data to the Olivine DER.
    - Rainforest Automation was unable to reproduce it on their own, and without remote access to the PG&E lab, this issue was not pursued.

  Note that in subsequent field testing with the newer Eagle 200 model, this issue has not been observed. One might infer this could point to the lab environment as being the culprit; however, a lab-tested Eagle was repurposed for field testing joined with a customer's SmartMeter™ and it continued to demonstrated this issue.

It is worth noting that the Eagle in use in the lab study relied on a Rainforest Automation proprietary API to send data to Olivine DER. This is not a particular issue for the lab nor for a limited field study where the objectives surround the ability of devices to provide telemetry from a SmartMeter™ to a cloud system capable of meeting the CAISO RIG requirements. In a wider production deployment, relying on a standard like OpenADR 2.0b would be recommended to reduce the chance of deploying assets that would be stranded if an alternate cloud vendor was required. Rainforest has indicated that OpenADR certification is in their development pipeline.

2.6.3 Universal Devices ISY-994 Challenges

There were no particular technical challenges raised with the ISY-994 devices; however, the following issues were noted during setup:

- The devices need to be specially ordered with the OpenADR VEN device certificate installed.
- The devices will attempt to communicate without encryption in the case that the VTN URL is entered without the HTTPS scheme prefix.
- The devices are highly configurable, and as such, are complicated to configure (noting that they also provide a high level of debugging information greatly easing onboarding in this project).

3 Findings

The following are findings from the lab study.

3.1.1 The CAISO 4-second requirement is not well supported by the configurations tested

The 4-second requirement is relevant for devices that are supporting spinning reserves. As identified in Section 2.1.1.5, the SmartMeter™ / HAN combinations utilized in the lab study have limitations that make these configurations unable to reliably provide 4-second telemetry.

3.1.2 Both HAN gateways are capable of providing 1-minute and 5-minute CAISO telemetry

As identified in Section 2.5.1, both of the HAN gateway devices are capable of meeting the 1-minute and 5-minute CAISO telemetry requirements. The 1-minute requirement is suitable for both energy and non-spinning reserves markets, while the 5-minute requirement is suitable only to energy.

3.1.3 The 1-minute instantaneous option has some benefits over the 5-minute option

As identified in Section 2.2, determining average demand over a relatively short interval from the SmartMeter™ using the HAN offers some challenges. While the average accuracy of such measurement methods should be equal to the average accuracy of the meter, the instantaneous reads are by their very nature as accurate as the meter can provide whenever provided. As such – and in combination with the broader applicability to market services – 1-minute telemetry may be the preferred frequency from the SmartMeter™ through a HAN gateway.

3.1.4 Off the shelf devices require some configuration to securely connect to cloud systems

In general, any device installed needs to be correctly configured to access a specific cloud system. This can be a default set by the manufacturer, a configuration added at deployment time by a third party, or configured by the customer. For example, "out of the box", the Rainforest Automation Eagle is configured to connect with the Rainforest cloud. For the purposes of the lab study, Rainforest had to add configuration and certificates specific to the project. In the Universal Devices case, the manufacturer installed the standard OpenADR client certificates, while the PG&E lab made the final OpenADR configuration changes necessary to connect with the Olivine VTN. In practice any wider deployment would rely on the former model, where the manufacturer would preconfigure the device in advance of shipping to customers. As such, a requirement for the selection of equipment should be the ability to create such pre-configured batches for shipment.

Note that this finding is not specific to HAN devices or the Olivine DER, but is part of the very nature of coordinating any client device with a cloud system.

3.1.5 Missing and divergent data occurred with both devices in the tested solution

There were two issues that were tracked through the project: data that did not arrive within a minute of the frequency cut-off (i.e., missing), and data that came either early or late (i.e., divergence). Both devices demonstrated these two cases of missing data and divergence from the scheduled frequency. In addition, both devices provided early data packets which was unexpected and warrants further investigation. The Rainforest Eagle had an additional issue that was related to missing packets – as covered in Section 2.4; however, this issue appears to be resolved in the newest model devices. All of these cases could undergo further evaluation to understand and perhaps resolve the underlying issues.

Note that the results for both devices followed a similar trend of less divergence and higher accuracy at slower polling frequencies. The relatively low rate of issues at the 1-minute and 5-minute frequency is an important basis for the finding that these devices can meet the CAISO requirements for energy and non-spinning reserves.

3.1.6 Potential Improvements to CAISO Direct Telemetry Rules

While the 2016 CAISO Direct Telemetry Business Practice Manual changes lowered the market barriers to providing telemetry from aggregated resources, this project identified some potential improvements that can be submitted to the CAISO through their BPM change process. The particular improvements in 2016 were directed at supporting existing investments in 5-minute real-time data collection activities, particularly around 5-minute average demand that was being collected from meter pulse outputs. While there are parallels with the average demand methodology used in the lab study, pulse outputs have unique issues where accuracy gets very poor at higher frequencies and / or low load levels. In addition, such meters do not allow an instantaneous read at all. With the experience of the HAN lab study, the following improvements could be made to simplify the requirements without sacrificing data quality and accuracy for SmartMeter™ to HAN gateway implementations:

- The 5-minute case requires a +/- 30 second clock synchronization between the devices. While clock synchronization in itself is not a concern since gateway devices can synchronize with meter time – an option unavailable to pulse devices – this clock alignment results in a 5-minute spike of traffic to the cloud provider. As the number of devices involved becomes very large in the mass market case, there is a clear value in distributing the transactions though time.
- The BPM considers 3-cases: 4-second, 1-minute, and 5-minute. Due to options with such gateway devices and other needs for this data, it would be helpful to codify other frequencies. For example, if 30-second operational telemetry would be useful to a product, it is presumably acceptable to serve the 1-minute instantaneous requirement. What is less clear is whether one could use 2.5 minute telemetry to serve the 5-minute requirement. This could be helpful to be resilient to an occasional drop of a value, allowing the 5-minute requirement to still be met.

In addition, one might infer from the BPM a conflict between the diagram in section 6.2.2 and the text in 6.2.3 in that the diagram identifies “average demand over 5 minute”, but the text clearly states that instantaneous demand is allowed in this case.

It is recommended that these suggested changes be recommended to the CAISO through their Business Process Change Management Process.

3.2 Future Research Questions

Several future research questions are raised by this lab study:

- How would a hybrid approach of different device types and cloud architectures (e.g., Embertek’s cloud solution paired with Rainforest Eagle direct push) perform compared to the existing solution architecture?
- The lab study was focused mostly on meeting the 1-minute and 5-minute CAISO telemetry requirements. A future research question could be what modifications to SmartMeter™ and/or HAN devices might be necessary to support the 4-second telemetry requirement.
- To support the spinning reserves without 4-second telemetry from the HAN, a future research question is how to mix 4-second telemetry with longer frequencies, particularly if a subset of the assets in an aggregation provide the frequency response requirement of spinning reserves. In that case, only the frequency responsive assets would need to provide 4-second telemetry.
- While the lab study focused on the delivery of CAISO telemetry, the real-time could also be utilized for distribution needs in addition – or instead of – CAISO telemetry. While standards for such data at the distribution level have not been established, how could these configurations feed into future standards work or pilot projects?

- The lab study focused on the residential SmartMeter™ case; however, many aspects of it are also relevant to Commercial cases. Studying the differences and similarities in the Commercial case could leverage the work of this project and other PG&E projects (e.g., the EPIC-funded commercial HAN project).
- The lab study noted some behaviors of the SmartMeters™ that could warrant further investigation. For example:
  - instantaneous readings did not appear to vary more frequently than every 16 seconds. Is this the behavior of the meter or indicative of an issue with the HAN devices. What is the highest frequency at which such data can be accurately retrieved? Similarly for summation data, such data did not appear to vary more frequently than once per 30 seconds, and also depended on the load. What are the explicit behaviors of the SmartMeters™ and can they be relied upon in the design of telemetry solutions?
- The general consensus of the 2015 DR Integration Working Group was that average demand over an interval is more appropriate at lower data collection frequencies than instantaneous demand readings. It may be worthwhile to test this consensus with analysis in the future.

3.3 Conclusions

The lab study sought to understand the behavior of two different HAN gateways under various scenarios, and ultimately, to measure the ability of each to support 1-minute and 5-minute telemetry options in support of CAISO telemetry. It was clear from the results that such devices are capable of supporting those requirements, noting that in the lab setting, the project team was insulated from the challenges and variability of on-premises deployment, including access to production SmartMeters™ and customer Internet access. As such, testing such devices in the field would uncover a greater level of detail on the costs of installation and support. Such a field study should take into account some of the questions raised in terms of data quality to ensure a greater clarity of reliability and accuracy can be determined.

Appendix A Example 4-second data

The following table shows a snapshot of 4-second payloads delivered by the Rainforest Eagle on 10/31/2016.

Reading time	Instantaneous kW
9:57:20 AM	0.32
9:57:24 AM	0.32
9:57:28 AM	0.32
9:57:32 AM	0.32
9:57:36 AM	0.319
9:57:40 AM	0.319
9:57:44 AM	0.319
*9:57:51 AM	0.319
9:57:54 AM	0.319
9:57:57 AM	0.319
9:58:00 AM	0.319
9:58:04 AM	0.319
*9:58:16 AM	0.318
9:58:20 AM	0.318
9:58:24 AM	0.318
9:58:29 AM	0.318
9:58:33 AM	0.571
9:58:37 AM	0.571
9:58:40 AM	0.571
9:58:44 AM	0.571
*9:58:54 AM	0.568
9:58:57 AM	0.568
9:59:00 AM	0.568
9:59:04 AM	0.568
9:59:08 AM	0.568
9:59:12 AM	0.568
9:59:16 AM	0.568
9:59:20 AM	0.566
9:59:25 AM	0.566
9:59:28 AM	0.566
9:59:32 AM	0.567
9:59:36 AM	0.567
9:59:42 AM	0.567
*9:59:52 AM	0.566
9:59:55 AM	0.566
9:59:58 AM	0.566

*Represents that at least one interval was not correctly sent to Olivine DER.

Appendix B Data Collection Testing Periods

The Eagle was tested during the following times. Note that the load banks were cycling between the two load profiles identified in Section 2.1.2 every hour.

Date	Start	Stop	Frequency (seconds)
July 20	21:00	23:00	20
July 21	06:00	08:00	10
July 27	00:00	08:00	10
July 27	12:00	14:00	60
July 27	16:00	18:00	150
July 27	19:00	21:00	300
August 1	15:00	17:00	300
August 1	19:00	21:00	150
August 1	22:00	23:00	60
August 9	09:00	11:00	60
August 9	12	14:00	150
August 9	22:00	23:00	60
August 9	15:00	17:00	150
August 9	19:00	21:00	300
August 9	23:00	1:00	10
August 10	8:00	10:00	20
August 10	12:00	14:00	30
August 10	15:00	17:00	60
October 21	8:00	10:00	20
October 21	12:00	14:00	30
October 21	15:00	17:00	60
October 24 - 31			300

The ISY-994 was tested during the following times:

Date	Start	Stop	Frequency (seconds)
August 18	08:00	10:00	10
August 18	11:00	13:00	20
August 18	14:00	16:00	30
August 18	17:00	19:00	60
August 18	20:00	22:00	150
August 18	23:00	01:00	300
October 20 - 23			300

Appendix C CAISO Telemetry Fit

This section describes the timing and latency requirements for PDR telemetry at the CAISO utilizing various measurement methods.

From the timing perspective then, the key aspects relevant to the lab test are:

A) The HAN polling frequency
B) Clock synchronization of the HAN device to a resource-specific synchronization time
C) The latency between the collection of the interval and the time at which the aggregate is available within the RIG for transmission to the ISO EMS.

These three aspects are covered in each of the sections below.

Note that the CAISO BPM allows for the election of either instantaneous demand or average demand over an interval. The following sections omit the 5-minute instantaneous demand option; however, the BPM does not prescribe the use of instantaneous or average demand readings tied to the frequency of the telemetry.

A review of telemetry requirements from the CAISO provides three levels of telemetry requirement based on the wholesale grid services being provided by a PDR resource[13]:

Grid Service	Frequency
Energy-only when PDR ≥ 10 MW	5-minutes or 1-minute
Non-spinning reserves	1-minute
Spinning reserves[14]	4-second

The 5-minute option was a result of the efforts of the Supply DR Integration Working Group completed in February 2015 and was added to the CAISO Direct Telemetry Business Practice Manual (BPM) in March 2016. Before that time, there was little specificity in the BPM related to the timing aspects of 1-minute telemetry.

5-Minute Telemetry (Average Demand Method)

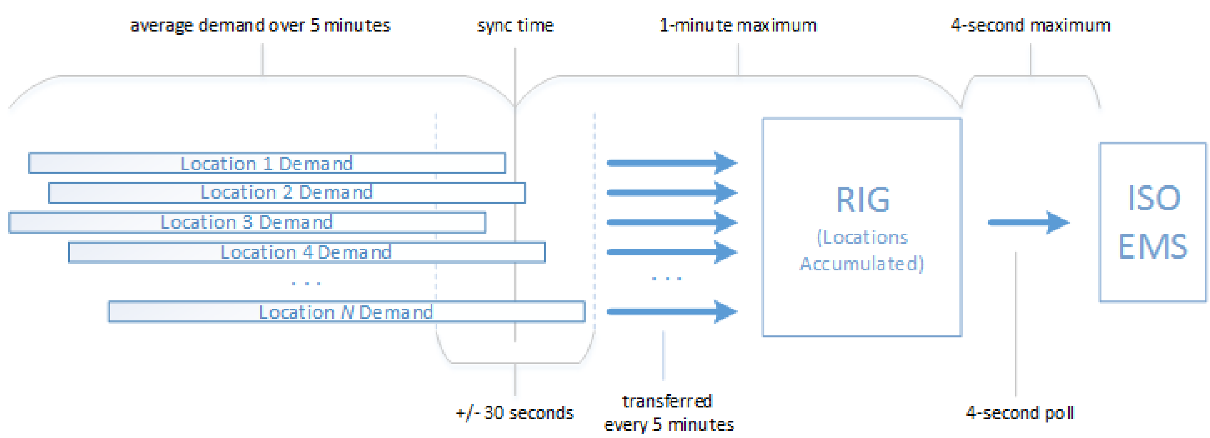

Figure 7: 5-Minute Telemetry for PDR[15]

[13] Frequency regulation is not included because it is not an approved grid service for PDRs at this time; however, if and when it is approved it is reasonable to expect it would have a 4-second telemetry requirement.
[14] Spinning reserves requires 4-second telemetry to validate compliance with its 8-second frequency response requirement.

Attribute	Notes
Polling Frequency	5-minutes.
Sync Time	When correctly configured, these systems stay synchronized to the meter time which keeps time within sync to +/- X, updated Y, meeting the +/- 30 seconds requirement.
Latency	The latency determined during testing was below 1 second, meeting the 1-minute maximum requirement.

1-Minute Telemetry (Average Demand Method)

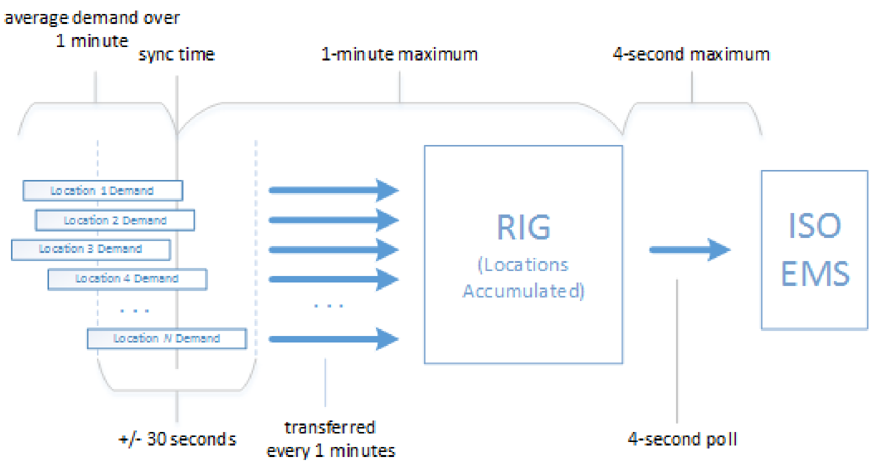

Figure 8: 1-Minute Telemetry for PDR[16]

Attribute	Notes
Polling Frequency	1-minute. Because of the 1-minute demand requirement – and barring the transmission of rolling averages – an increased frequency does not apply.
Sync Time	The same comment for 5-minute frequency applies here, noting that a +/-30 second requirement for 1-minute reads makes the actual clock synchronization state irrelevant.
Latency	The latency determined during testing was below 1 second, meeting the 1-minute maximum requirement.

[15] From the CAISO Direct Telemetry BPM version 9. See https://www.caiso.com/rules/Pages/BusinessPracticeManuals
[16] While not a part of the CAISO BPM, this diagram is implied by the textual description of 1-minute telemetry in the BPM.

1-Minute Telemetry (Instantaneous Demand)

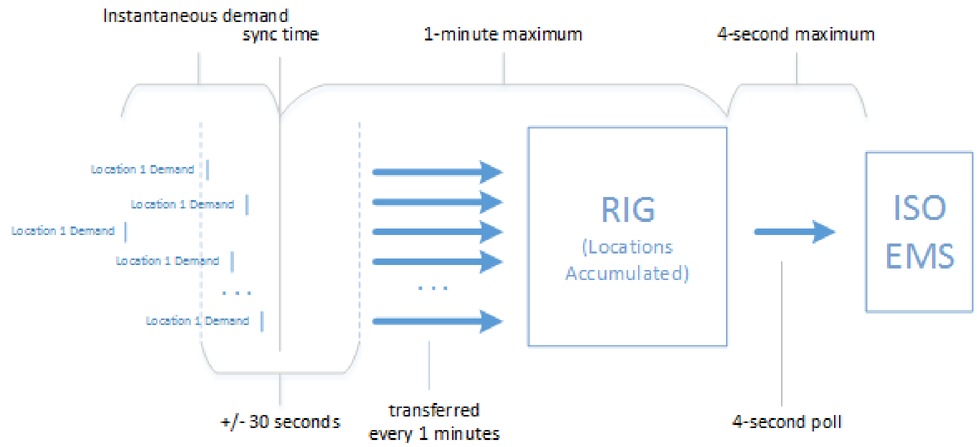

Figure 9: 1-Minute Telemetry for PDR (Instantaneous)

Attribute	Notes
Polling Frequency	1-minute. An increased frequency could ameliorate issues of divergent data.
Sync Time	Again in this case, the state of synchronization is irrelevant.
Latency	The latency determined during testing was below 1 second, meeting the 1-minute maximum requirement.

Appendix D Data Accuracy Proposals

The following text comes from the final report of the 2016 Supply DR Integration Working Group sub-working group on Telemetry. This section is particularly related to data accuracy and contemplates the specific concerns about how the CAISO accuracy requirement impacts aggregations. This issue and proposal were generally agreed to by the CAISO at the time; however, the output of the working group has not been incorporated into the CAISO Business Practice Manual (BPM). It is Olivine's position that an abbreviated form of this proposal should be added to the BPM, noting that it may result in additional compliance requirements for DRPs.

Issue Statement

Section 5.5, Data Validation and Confidentiality of the Direct Telemetry BPM states the requirement for telemetry accuracy:

> *All telemetry data reported via the RIG must be within +/-2% of the true value.*
>
> *The CAISO or its designee may inspect the resource owner's RIG and related facilities to verify the accuracy and validity of all data telemetry to the CAISO. The CAISO reserves the right to periodically audit and re-verify the accuracy and validity of all telemetry data. In addition, the CAISO's verification activities will be coordinated with the resource owner at least 24 hours in advance.*
>
> *. . .*
>
> *All data telemetry provided through the resource owner's RIG shall be tested by the resource owner or resource owner's representative for accuracy and validity on a periodic basis as necessary to assure that the accuracy requirements are maintained. The best practice is to test all resource data annually for accuracy.*

For actual "direct telemetry" – as is the case for conventional generators – the +/-2% accuracy requirement is mainly one of ensuring adequate equipment (i.e., metering, current transformers, etc.) that is capable of providing an instantaneous read within +/-2% of true power. The measurements from such equipment are provided directly to an on-site RIG from which it is collected by the CAISO.

When applying this to aggregations of locations within a PDR, there are various concerns:

- When measurements are performed using a KYZ pulse meter, it may be that suitable accuracy can only be achieved during certain times of the day. This is because the accuracy of this technology within any interval of time is related to the number of pulses generated during that time. So, for example, a specific location may be calibrated to provide KYZ pulses that meet the accuracy requirement at intervals shorter than the utility metering interval, but only during periods of sufficient load. It may be less accurate when a PDR is dispatched, because the KYZ pulse meter produces fewer pulse at lower loads[17]. Thus, accuracy must be determined under different load levels.
- When measurements are performed using either instantaneous reads or average interval reads that are inherently accurate across the interval, suitable accuracy can be achieved as long as a sufficient number of locations are correctly reporting data.
- Unlike meter data provided for billing and settlements, there is no opportunity to perform Validation, Editing, and Estimation (VEE) on real-time telemetry. As a result, the telemetry data are "raw" and may at times include gaps and spikes that would normally be corrected in a normal VEE process. This is the same problem that conventional generators have (i.e., that raw data is provided), but it is made more challenging because of the number of locations that are combined to produce telemetry for PDR.

[17] The KYZ technology has an inherent dead band within which the meter will not produce any pulse counts. This dead band is more of a concern as the interval period becomes shorter.

It is the belief of the parties that the telemetry accuracy requirement must be resolved in a cost-effective manner, or it may inhibit the ability of PDR to participate in the ancillary services markets and possibly the energy markets at a scale large enough to support a valid business model.

Proposed Solution

The parties would like the CAISO to accept the following methodology for meeting the requirement:

- The resource owner will test at installation time that individual location-installed telemetry solutions are within +/-2% accuracy compared against test intervals using the billing meter or using an instantaneous energy measurement device separate from the telemetry solution that is within +/- 2% accuracy, noting that in the case of comparing against billing intervals, an entire billing interval will be used for comparison.
- For KYZ and any technology for which the accuracy may be dependent on interval and load level, the resource owner will test data accuracy under different conditions, including normal load conditions and during the period when the PDR is dispatched. It will be important to demonstrate that the data meet accuracy requirements at periods of low load as in a dispatch situation.
- The resource owner will produce an annual telemetry report (see below for an example specific to KYZ metering) that demonstrates accuracy compliance. This report will also include reports on PDR location "uptime" to validate accuracy for non-KYZ metering.
- The CAISO is asked to acknowledge and/or make an exception explicitly allowing that telemetry data is "raw" and may include short-term anomalies. It will continue to be the resource owner's responsibility to correct any persistent issues.

The resource owner will be responsible for ensuring that the +/-2% accuracy will be met during periods of market activity (i.e., when bids are submitted). This would specifically allow KYZ metering that is calibrated for periods where load curtailment would be available to the market.