THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE

The business case for behind-the-meter energy storage

Q1 performance of UQ's 1.1MW Tesla battery

Contents

UQ Project Team

Andrew Wilson
Senior Manager – Energy & Sustainability
a.wilson@pf.uq.edu.au

Danielle Esterhuysen
Program Manager – Energy
d.esterhuysen@uq.edu.au

Dominic Hains
Energy Engineering Officer
d.hains@pf.uq.edu.au

Installation and commissioning of the battery was provided by CSA Services and QGE Group, with technical specifications developed by GHD.

The project team would like to thank Enel X for their assistance with preparing this report.

Note that all values in this report are in $AUD and are exclusive of GST unless otherwise stated.

Published: May 2020

1. Executive Summary

As part of the organisation's energy leadership ambitions, The University of Queensland installed the state's largest behind-the-meter battery in late 2019. The 1.1MW / 2.15MWh Tesla Powerpack system accompanied UQ's move to be the first university in Australia to participate directly in the wholesale electricity spot market. At an all-in cost of $2.05 million, the project was funded through the sale of renewable energy certificates created by UQ's existing 6.3MW behind-the-meter solar PV portfolio.

This report explores the performance of the battery in-depth during Q1 2020, including its revenue, a comparison to business case assumptions, the effectiveness of its control strategy, technical issues and challenges, and key learnings. It aims to provide a transparent 'warts and all' look at the opportunities and challenges of utilising behind-the-meter battery storage to generate revenue and reduce energy costs.

Table 1.1: Battery key figures

Make & model	Tesla Powerpack 2.5
Rated power	1.11 MW
Storage capacity	2.15 MWh (~2 hours at full power)
Depth of discharge	100% of nameplate
Number of battery packs	10 x 215 kWh
Number of inverters	2 x 580 kVA (at 415 Volts)
Physical footprint	44 m² (including clearances)
Total weight	25.7 tonnes (excluding foundation)
Total project cost	$2.05 million ($954/kWh)
Date commissioned	19 November 2019

Revenue Streams

The UQ battery has been developed to deliver revenue and value from the combination of four distinct services:

Arbitrage

A custom developed control system (the 'Demand Response Engine' or DRE) aims to charge the battery when prices are low and discharge when prices are high - maximising the spread between prices to help offset energy costs while respecting the fact that the battery only has a finite storage capacity (roughly two hours at full power). Refer to section 3.2 for further information.

Peak Demand Lopping

It is intended that the battery will help UQ to reduce its monthly peak demand charges by lopping the top off the highest demand intervals of each month. The control strategy that enables this functionality was not completed in time for Q1 2020. Refer to section 3.5 for further information.

Contingency FCAS

Through a partnership with Enel X, the battery is paid to remain on standby to respond to sudden disturbances to grid frequency from events such as power plants tripping offline. Revenue is earnt by bidding this response capability into the NEM's three contingency Frequency Control Ancillary Services (FCAS) markets. Refer to section 3.3 for further information.

Virtual Cap Contract

As a spot price exposed customer, UQ is required to put hedging strategies in place to prudently manage risk. One option available is the use of cap contracts which limit financial exposure to extreme prices (typically >$300/MWh). These hedging products can be considered as a form of insurance. The battery is able to provide this insurance 'virtually' in place of buying a cap contract by responding quickly to high price events and minimising UQ's exposure. Whilst not an exact replacement for a traditional financial cap, this service has substantial value to UQ nonetheless. Refer to section 3.4 for further information.

Q1 Performance

Across the three months of Q1 2020, the battery delivered a total of **$74,000** in revenue across its four services. This is broken down in Figure 1.1. This total exceeded business case assumptions by over 20%, despite no value being derived from peak lopping during the quarter. These results were primarily driven by a 54% overperformance in contingency FCAS revenue compared to forecast, largely due to unprecedented network conditions across the NEM related to the Black Summer bushfires combined with storm events.

Figure 1.1: Breakdown of Q1 revenue by value stream

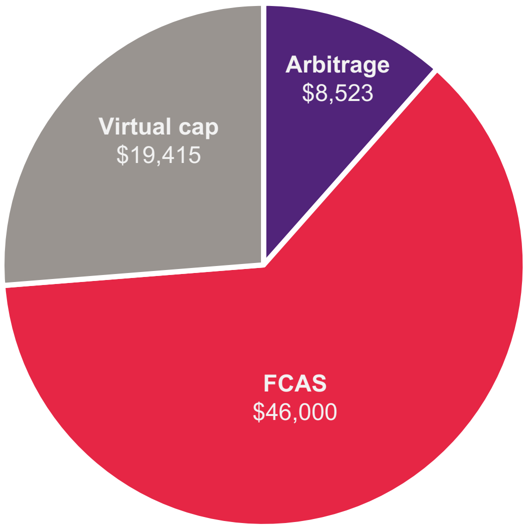

Acknowledging that these figures are representative of one quarter only - and that past performance is not an indicator of future results - they are nonetheless a promising indication that the battery remains on track to meet or exceed its forecast 8-year payback. It is reasonable to expect that future performance will see the breakdown between revenue streams differ from quarter to quarter, and that the value of some services may diminish over time (e.g. contingency FCAS), whilst others could be expected to escalate (e.g. arbitrage) as the energy transition across the NEM evolves. Furthermore, these figures do not account for the potential of future new revenue streams, such as participation in the regulation FCAS market, or emerging fast frequency response and virtual inertia markets.

Issues & Learnings

Despite performance across Q1 being successful overall, a number of key issues and learnings emerged. These primarily related to the challenges of developing an effective control strategy to maximise arbitrage revenue. These efforts are hindered by the inherent unreliability of the AEMO pre-dispatch price forecasts upon which the control strategy currently relies. This is then compounded by issues of deciding between price certainty in the moment ("one in the hand") versus the potential for higher - but more uncertain - returns later ("two in the bush"). These issues and ideas for helping to address them are discussed further in sections 4 and 5.

The battery also spent a cumulative total of 124.5 hours (5.7%) of the quarter offline. This was primarily a result of nuisance tripping related to electrical protection settings that have since been tuned to avoid the issue. Problems were also encountered several times during the quarter with network outages impacting the control system architecture, causing the battery to revert to outdated charge/discharge patterns or requiring manual intervention. These issues and plans for future resiliency to help address them are also discussed further in sections 4 and 5.

Table 1.2: Q1 2020 key performance figures

Total MWh charged	107.70 MWh
Total MWh discharged	90.98 MWh
Round trip efficiency*	84.5%
Average charge price	$43.20/MWh
Average discharge price	$149.98/MWh
Average 'spread'	$106.77/MWh
Charge price compared to average spot price	-20%
Discharge price compared to average spot price	+178%
Battery availability/uptime	94.3%
Capacity factor (discharging only)	3.8%
Capacity factor (charging + discharging)	8.3%
Number of contingency FCAS events	12
Total contingency FCAS response duration	42 minutes 36 seconds

*Figure includes auxiliary load and is impacted marginally by the % capacity the battery started and ended the quarter with. This figure compares to 86.5% stated round-trip efficiency under nominal conditions, noting that Q1 includes the hottest months of year and higher than average auxiliary cooling load.

2. About the Battery

2.1 Organisational Drivers for Install

In October 2017, the UQ Senate approved the business case for the Warwick Solar Farm initiative and set UQ on the path to fundamentally change how the organisation consumes and procures electricity. This included approval to become the first university in Australia to participate directly in the wholesale electricity spot market. As a large energy generator and a large energy consumer (a 'Gensumer'), UQ now has the relatively unique ability to leverage the opportunities of being a participant on both sides of the energy market to maximise value and deliver UQ's energy needs in a flexible, sustainable, and lowest cost manner. The transition from a passive retail electricity customer to an active participant in the wholesale electricity market requires UQ to now control not just the *quantity* of how much electricity is being used, but also the *timing* of when this electricity is being consumed.

Demand Response (DR) and energy storage were identified as key pillars of UQ's Gensumer energy strategy in order to make use of abundant 'free' energy from the Warwick Solar Farm during the day, shape the campus energy demand profile, and to respond to electricity spot price market volatility in order to minimise price exposure.

The first step was to examine what DR capabilities UQ already had behind-the-meter, how to expand this capability, and how to autonomously control the DR portfolio to respond to spot price fluctuations, including outside of regular business hours. It became apparent that a megawatt-scale lithium ion battery system would be the most appropriate next step for supporting UQ's energy transition ambitions. Further scoping also determined that a custom-built control system was necessary both to ensure compliance with UQ's IT security requirements, as well as enabling maximum flexibility to integrate and control a variety of DR assets. Finally, as an education institution it was vital that the battery was installed in a prominent location at UQ's main St Lucia campus in order to ensure it was accessible to staff and students and could maximise additional value from teaching, research, and engagement opportunities.

UQ has previous experience with the installation and operation of energy storage behind-the-meter. A pilot 600 kW / 750 kWh lithium ion battery was installed in 2016 as part of the 3.3 MW Gatton Solar Farm, several zinc bromide batteries have been trialled since 2011, and a 150 kW / 600 kWh vanadium redox flow battery was recently installed at the Heron Island Research Station as part of an off-grid hybrid renewable power station.

The Tesla battery – known internally as the 'Engineering Precinct Battery' – is UQ's first installation undertaken primarily for commercial purposes. For this application, lithium was chosen for its energy density, technical capability to quickly swing between charge and discharge, ability to participate in a potential future Fast Frequency Response market, and maturity of product offerings. A key learning from previous battery installations was also to minimise complexity and interfaces wherever possible – off the shelf solutions with components that have been tested and certified to work with each other as a complete package are preferable and more cost effective in the long run.

2.2 Business Case

In mid-2018, large-scale generation certificates (LGCs) were trading close to their ceiling price, with it being widely forecast that the value of these certificates would fall in the near future as supply sharply increased. UQ subsequently made the decision to cash in its sizable holding of LGCs that had been generated from rooftop solar installations and the Gatton Solar Farm. This money was set aside to invest in the next frontier of energy initiatives for UQ – a re-investment of revenue from UQ's previous renewable energy leadership helping to enable further investment towards UQ's clean energy goals.

A business case for the battery was prepared in mid-2018 to ensure this funding was used on a project that would represent value-for-money to the University while also leveraging additional teaching, learning, and engagement opportunities. The battery's forecast financial return to UQ was calculated based on the assumption that several different value streams could be 'stacked' - arbitrage, peak demand lopping, FCAS, and replacement of financial cap contracts with a 'virtual' cap contract.

Wholesale electricity price forecasts were used to estimate potential arbitrage income, while backtest simulations were run to estimate potential savings from peak demand lopping. FCAS revenue estimates were derived from market forecasts and previous experience of the Gatton battery's participation in Enel X's FCAS aggregation scheme. Savings from avoided cap contracts were based on wholesale price forecasts and futures pricing from the ASX. The capital cost of the project was estimated at $2 million – derived from an assumed value of $1,000/kWh for a 2MWh project. The size of the project was largely determined by the amount of funding available, as well as constraints around available transformer capacity on site.

Based on this financial modelling, the business case estimated that the combination of these revenue streams from the battery would provide a financial return of around $245,000 per annum, with an estimated payback of 8 years. This estimate sits within the 10 year warranted life of the system and its expected technical life of 15 years.

2.3 Capital Cost & Construction

UQ engaged GHD to help develop a technical specification for the project, and a tender was issued to pre-qualified contractors in late 2018. CSA Services and QGE Group was ultimately selected to deliver the project for UQ utilising the Tesla Powerpack 2.5 battery product, which was selected for a range of commercial and technical reasons, as well as its well-established track record and integrated 'all in one' approach – a factor that was important to UQ as discussed above.

Difficulties were encountered with finding a suitable installation site that was above the flood plain, was nearby to sufficient transformer capacity, did not pose a fire risk indoors, and would not obstruct other functions of the University on an already crowded campus. A final location was chosen between two buildings in the south-east area of the campus where a disused garden bed was located. In early 2019, site prep works began to ready the area for install. Although the battery order was placed in December 2018, production of the Tesla Powerpack 2.5 did not start until mid-2019. This meant that the battery was not delivered to Australia until September 2019. Once on site, installation of the battery components was completed extremely fast, with battery packs being dropped in to place in the space of an afternoon. Accounting for site prep works and follow up wiring and commissioning, the total time to deliver the project was around 2 months.

The final 'all in' cost of the project, including consultants, EPC contract, and other ancillary works (but not accounting for internal project management costs) was $2.05 million or $954/kWh. A breakdown of this cost by component is provided in Table 2.1.

Visit https://tinyurl.com/y7jsoqk4
for more information

Table 2.1: Breakdown of battery capital cost

	Cost (ex GST)	$ / kWh
Battery supply cost*	$1,700,000	$791
Battery balance-of-plant and Comissioning	$182,000	$84
Site Prep & Construction	$135,000	$63
Soft Costs	$35,000	$16
Total Cost	$2,052,000	$954

*Third-party supplier cost as part of EPC contract. AUD/USD foreign exchange rate was $1.40 in Dec 2018.

Figure 2.1: Live Data Display

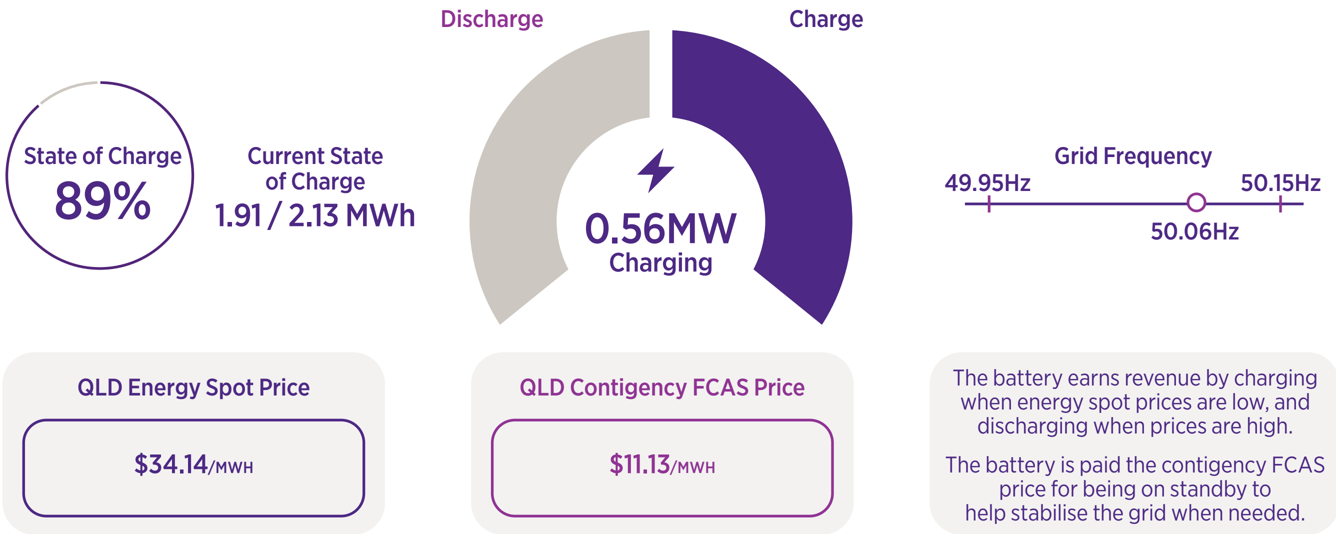

2.4 DRE Control Strategy

Given the current trends associated with the Internet of Things, as well as advancements in modern computing technologies, there is an ever-growing movement for software solutions to supplement and optimise traditional approaches to SCADA in the industrial automation industry. The overall performance of such a cloud service can be largely attributed to the effectiveness of data management within the system, namely:

- Data Collection,
- Data Standardization,
- Data Analysis,
- Data Evaluation and,
- Data Storage/ Retention.

Whilst it was necessary to review and consider trends and implementations across industry, it was also critical for UQ to design and evaluate such a system with business specific objectives in mind. By recognising system performance, system flexibility, and data availability as key requirements for the system to meet UQ's needs, the decision to build a platform-as-a-service (PaaS) in-house over purchasing an enterprise platform became a simple one. It was also determined that an in-house system would have similar up-front development costs to the establishment of an enterprise platform, but with significantly reduced ongoing licensing fees. Considering this, the concept of the Demand Response Engine – known as DRE – was born as a solution for UQ to manage not only the battery, but additional demand response initiatives, such as HVAC control, as they were developed.

DRE is a cloud-based, data-driven, supervisory control system hosted within Amazon Web Services. DRE is a novel platform in which autonomous, event-driven predictive controllers can be designed, simulated, and deployed across UQ infrastructure to help improve and optimise energy asset operation. Figure 2.2 illustrates the high-level architecture of the system.

To provide supervisory control for the battery, DRE utilises a model predictive control (MPC) approach. By utilising a system model for the behaviour of the battery, the scheme entails repeatedly solving a constrained cost optimisation problem, using predictions of future energy costs over a moving time horizon (i.e. as new forecast or pricing information is received) to choose an appropriate control action for the battery. Constraints defined for the optimisation problem include operational constraints of the physical battery system, energy constraints associated with FCAS commitments, and peak demand limiting constraints derived from site level power metering. The receding horizon control approach is possible for the financially driven control of the battery system because markets in the NEM are slow sampling, updating every 5-minutes. The output from DRE is a 30-minute resolution control policy of real power set points synchronised with trading periods in the NEM. This control policy is written to the battery controller via Modbus TCP. As new market information is received every 5 minutes, DRE recomputes the optimal policy for the battery, causing the control schedule to update.

Figure 2.2: High level overview of DRE architecture

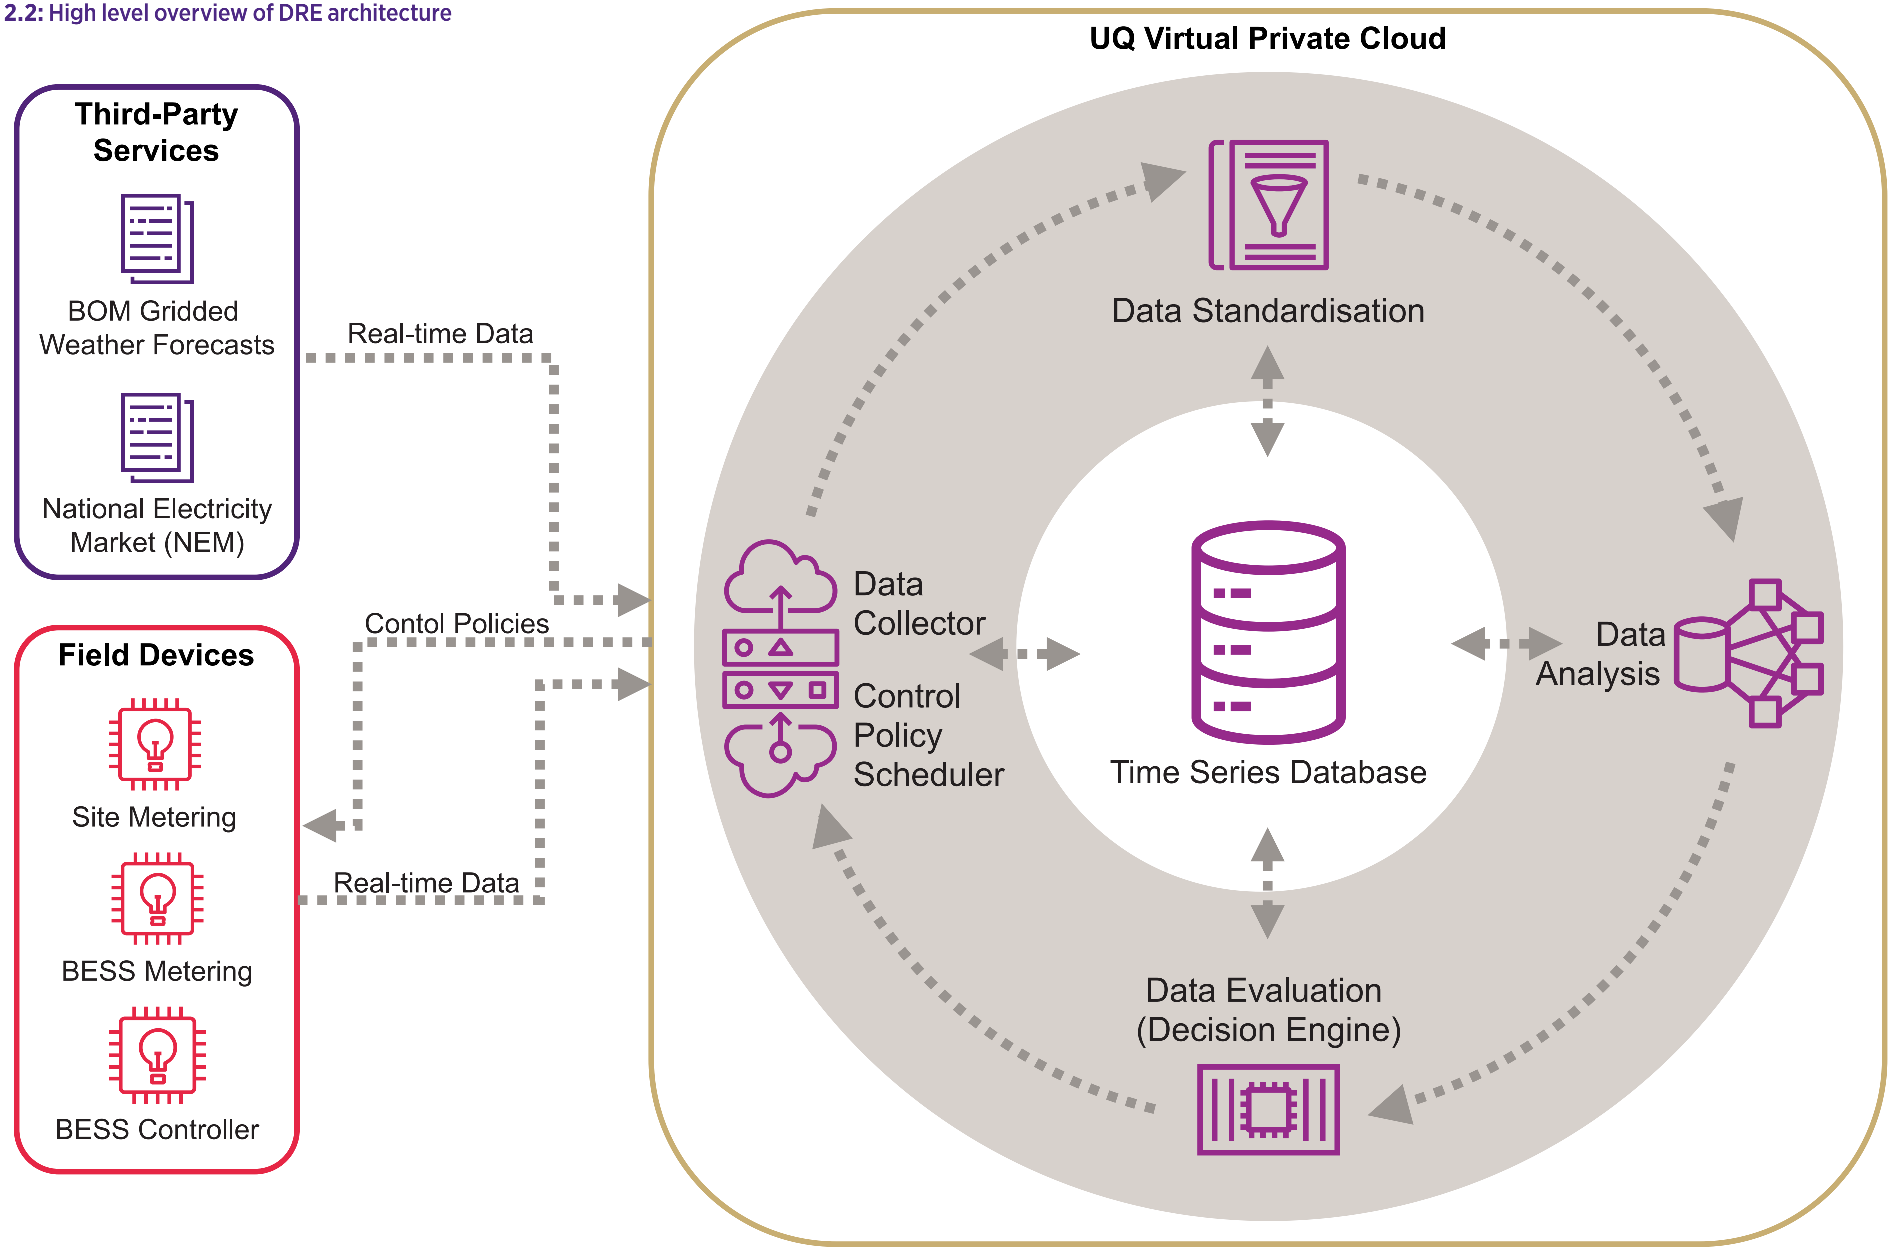


THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE
OFF
ON
TESLA
TESLA

3. Q1 2020 Performance

3.1 Overall Performance

In total, the battery delivered $73,938 in value during Q1 2020. This was dominated by FCAS, which delivered 62% of total revenue, followed by the virtual cap contract at 26%, and finally arbitrage at 12%. Each of these revenue streams is discussed in depth in the following sections. Total revenue in Q1 exceeded business case forecasts by just over 20%. Comparison of forecast and actual revenue by service is illustrated in Figure 3.1.

Overperformance of the total revenue forecast was driven primarily by FCAS, which exceeded forecasts by 53% for the reasons discussed in section 3.3. The virtual cap contract also overperformed forecasts, while arbitrage underperformed marginally for the reasons discussed in section 3.2. Notably, no revenue was earned from peak demand lopping, as this functionality was not ready in time for the start of Q1, as further discussed in section 3.5.

Figure 3.1: Comparison of Q1 actual versus forecast revenues by stream

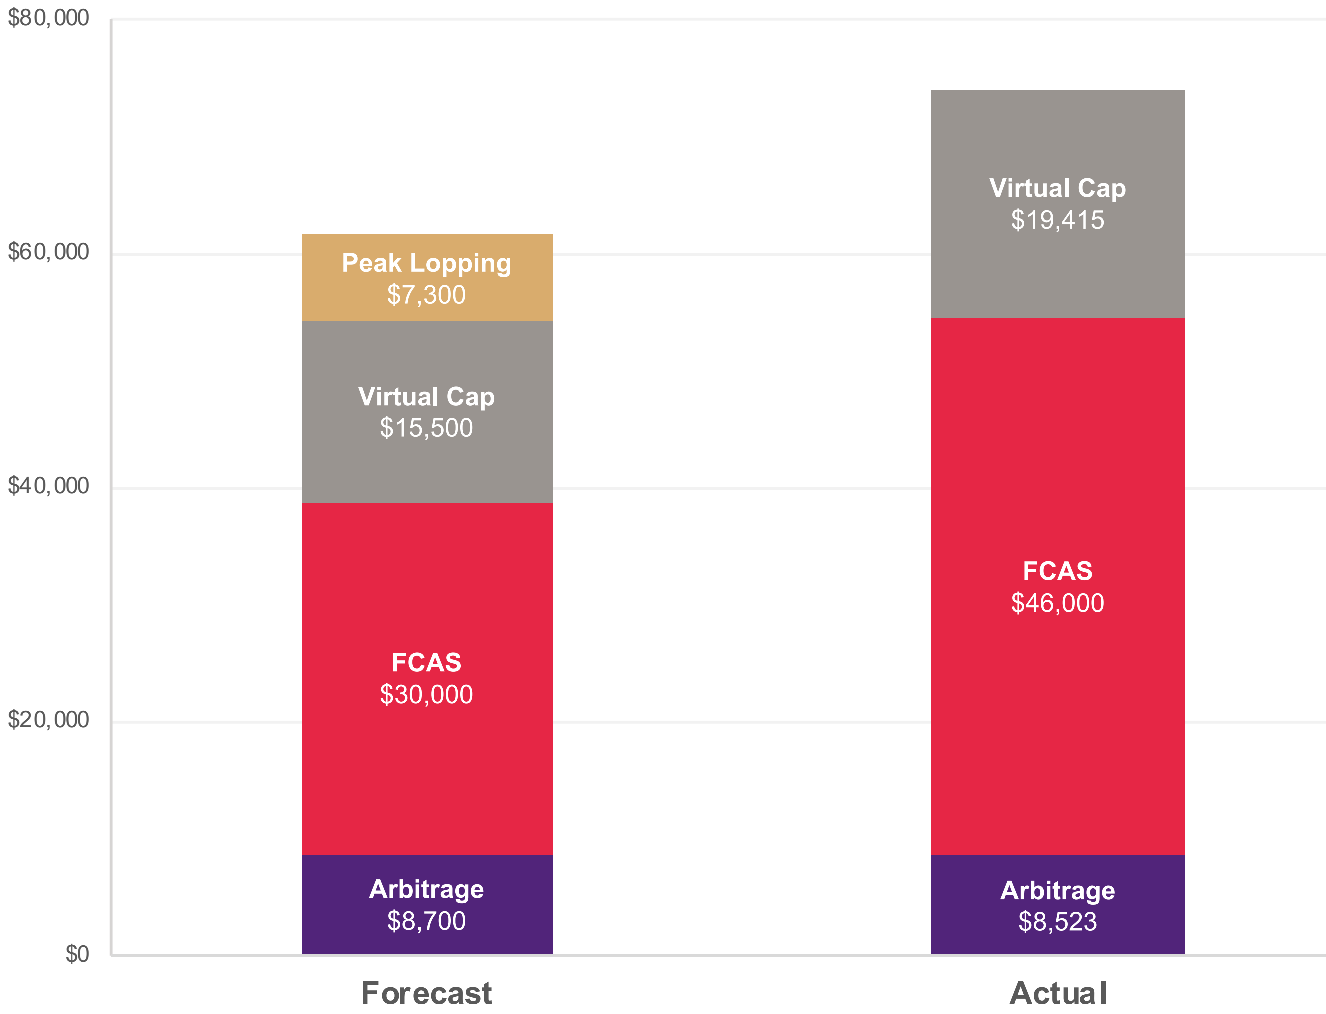

3.2 Arbitrage

As a spot price exposed energy user, one of the core functions of UQ's battery is to undertake arbitrage – charging to store energy when prices are low and discharging to generate energy when prices are high. While a simple concept, being able to effectively select the lowest and highest prices in a volatile market such as the NEM presents unique challenges. These are discussed in further depth in section 4, but include a reliance on imperfect forecasts, as well as what is known as the '5/30 rule' whereby spot prices are set every 5 minutes but financially settled based on a half hourly average.

UQ's battery utilises DRE's supervisory control system to automatically trade in the NEM 24 hours a day, 7 days per week. With the exception of one event (discussed in section 4.2), arbitrage performance during Q1 involved no manual intervention and is a result of DRE's algorithm-based trading only. DRE is configured to schedule battery charging and discharging to achieve the maximum spread across each forecast horizon of up to 40 hours into the future. This schedule is reviewed and revised every five minutes as AEMO's five minute and thirty minute look-ahead pricing forecasts are updated. This in turn allows DRE to schedule the battery without any need to configure a bias towards periods which would be typically considered 'peak' or 'off-peak' times.

DRE's ability to direct the battery is constrained by the need to achieve a minimum spread or else charging and/or discharging will not be scheduled to occur. This is required to ensure that the battery does not chase every spread that it sees in the look-ahead pricing forecasts. Determining the threshold for this minimum spread involves weighing up a trade-off between wanting to ensure that the battery's full volume is utilised and capacity factor maximised, but also ensuring that round-trip efficiencies are accounted for and that the implied cost of degradation is considered. These factors effectively create a floor to the minimum viable spread price below which it is uneconomic to undertake arbitrage. DRE's arbitrage controller is also constrained by the warranty conditions on the battery, with these effectively acting to limit overall throughput per annum to a level that equates to one full charge and discharge cycle per day.

Significant debate occurred regrading whether the reference charge price in the spread calculation should be retrospective (i.e. based on what the battery last 'filled up' at) or prospective (i.e. what the battery could fill up at in the future). Both options present pros and cons, and it was ultimately determined that a prospective approach was likely to result in the least instances of unintended behaviours. This is an area that is continuing to be monitored and improved, particularly in the context of the issues discussed further in section 4. The need to actively adjust the minimum spread threshold up and down based on forecast levels of spot pricing and volatility across different months and quarters is also an issue currently being further explored.

Figure 3.2 provides an example of the battery's performance on 4 January – the first day of high prices and volatility in the quarter. Figure 3.3 provides an example of DRE's arbitrage decision making in a different context – discharging during the evening of 14 March despite prices being relatively 'low' in order to free up space to charge the following day during forecast negative price intervals.

Figure 3.2: Battery performance during market volatility on 4 January

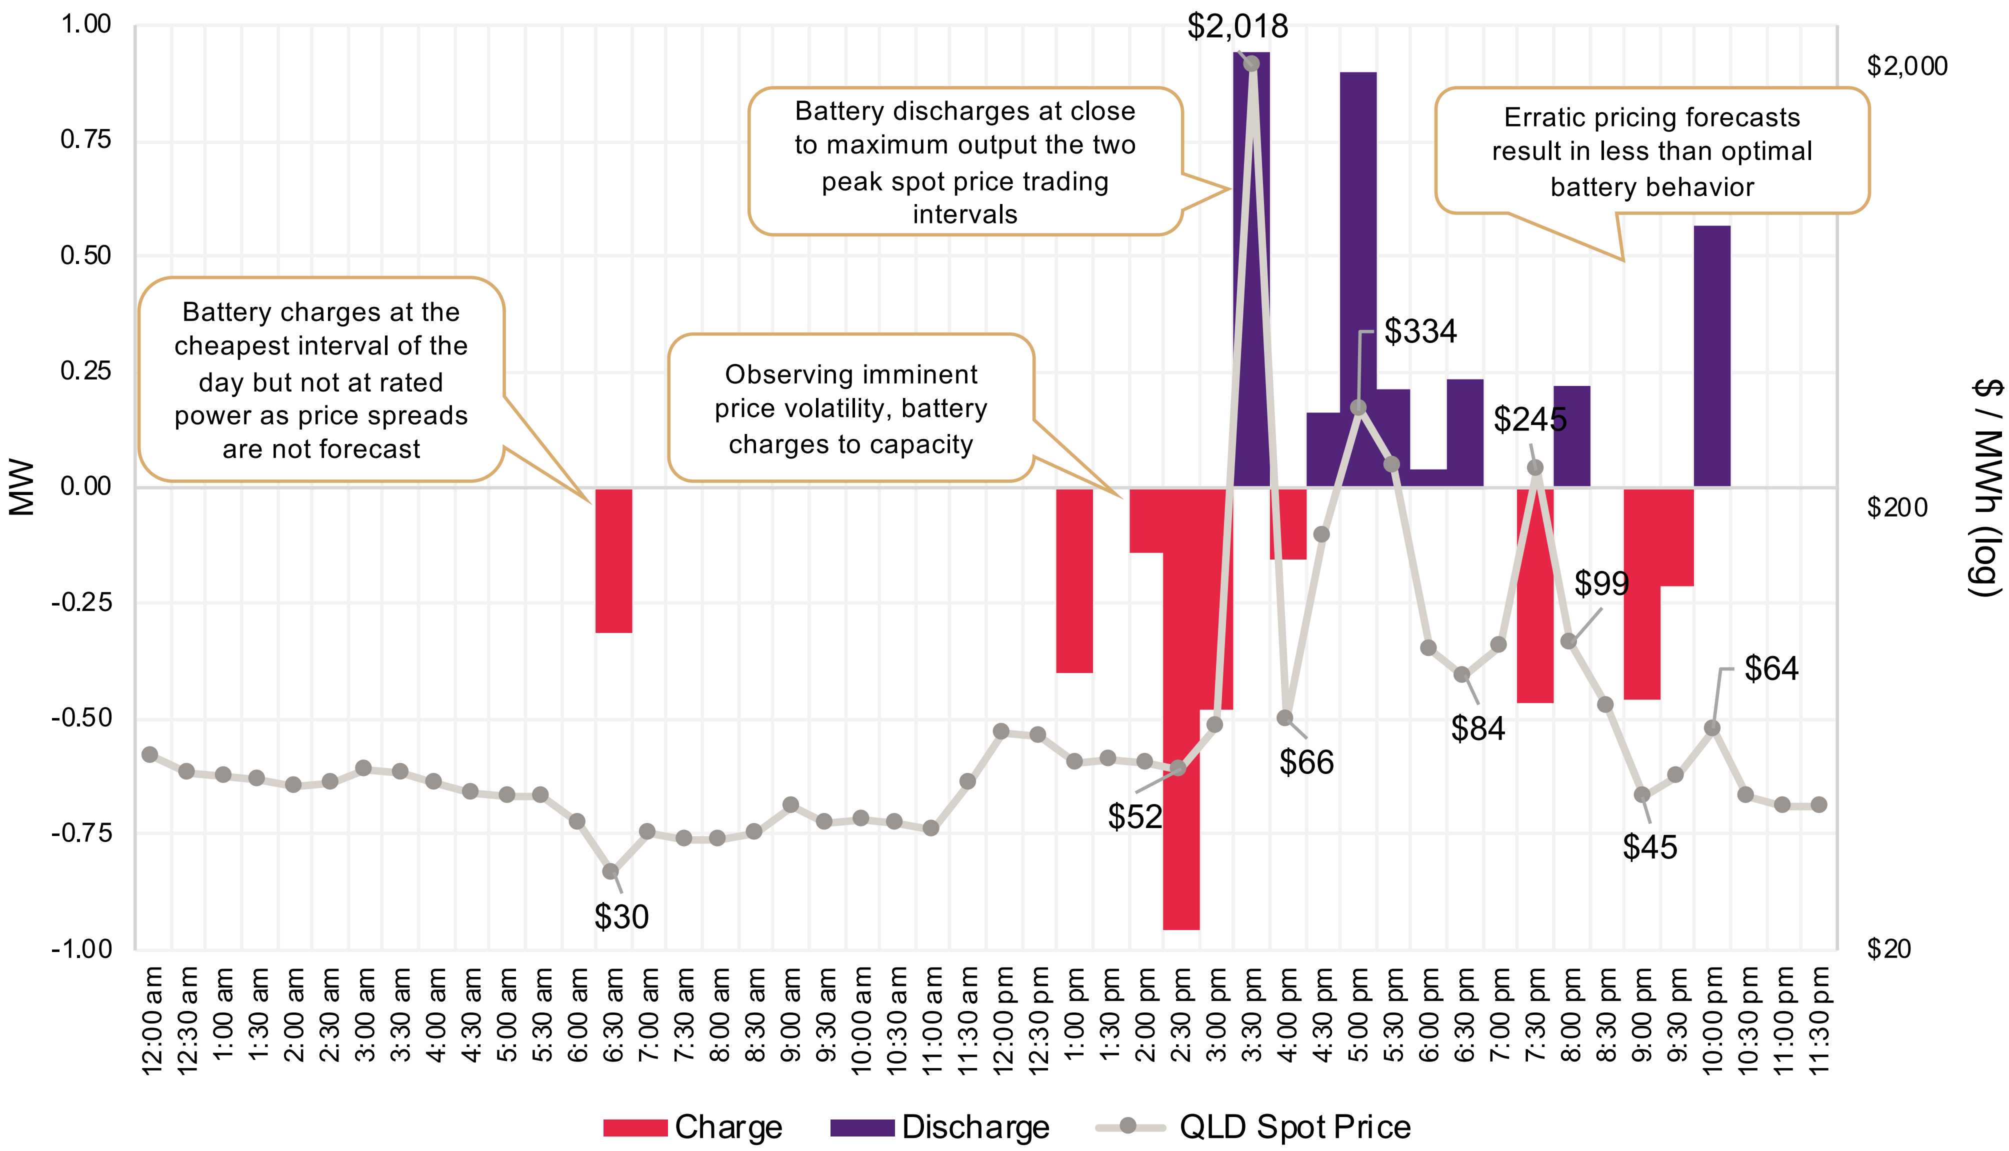

Figure 3.3: Battery performance across 14 + 15 March during negative pricing intervals

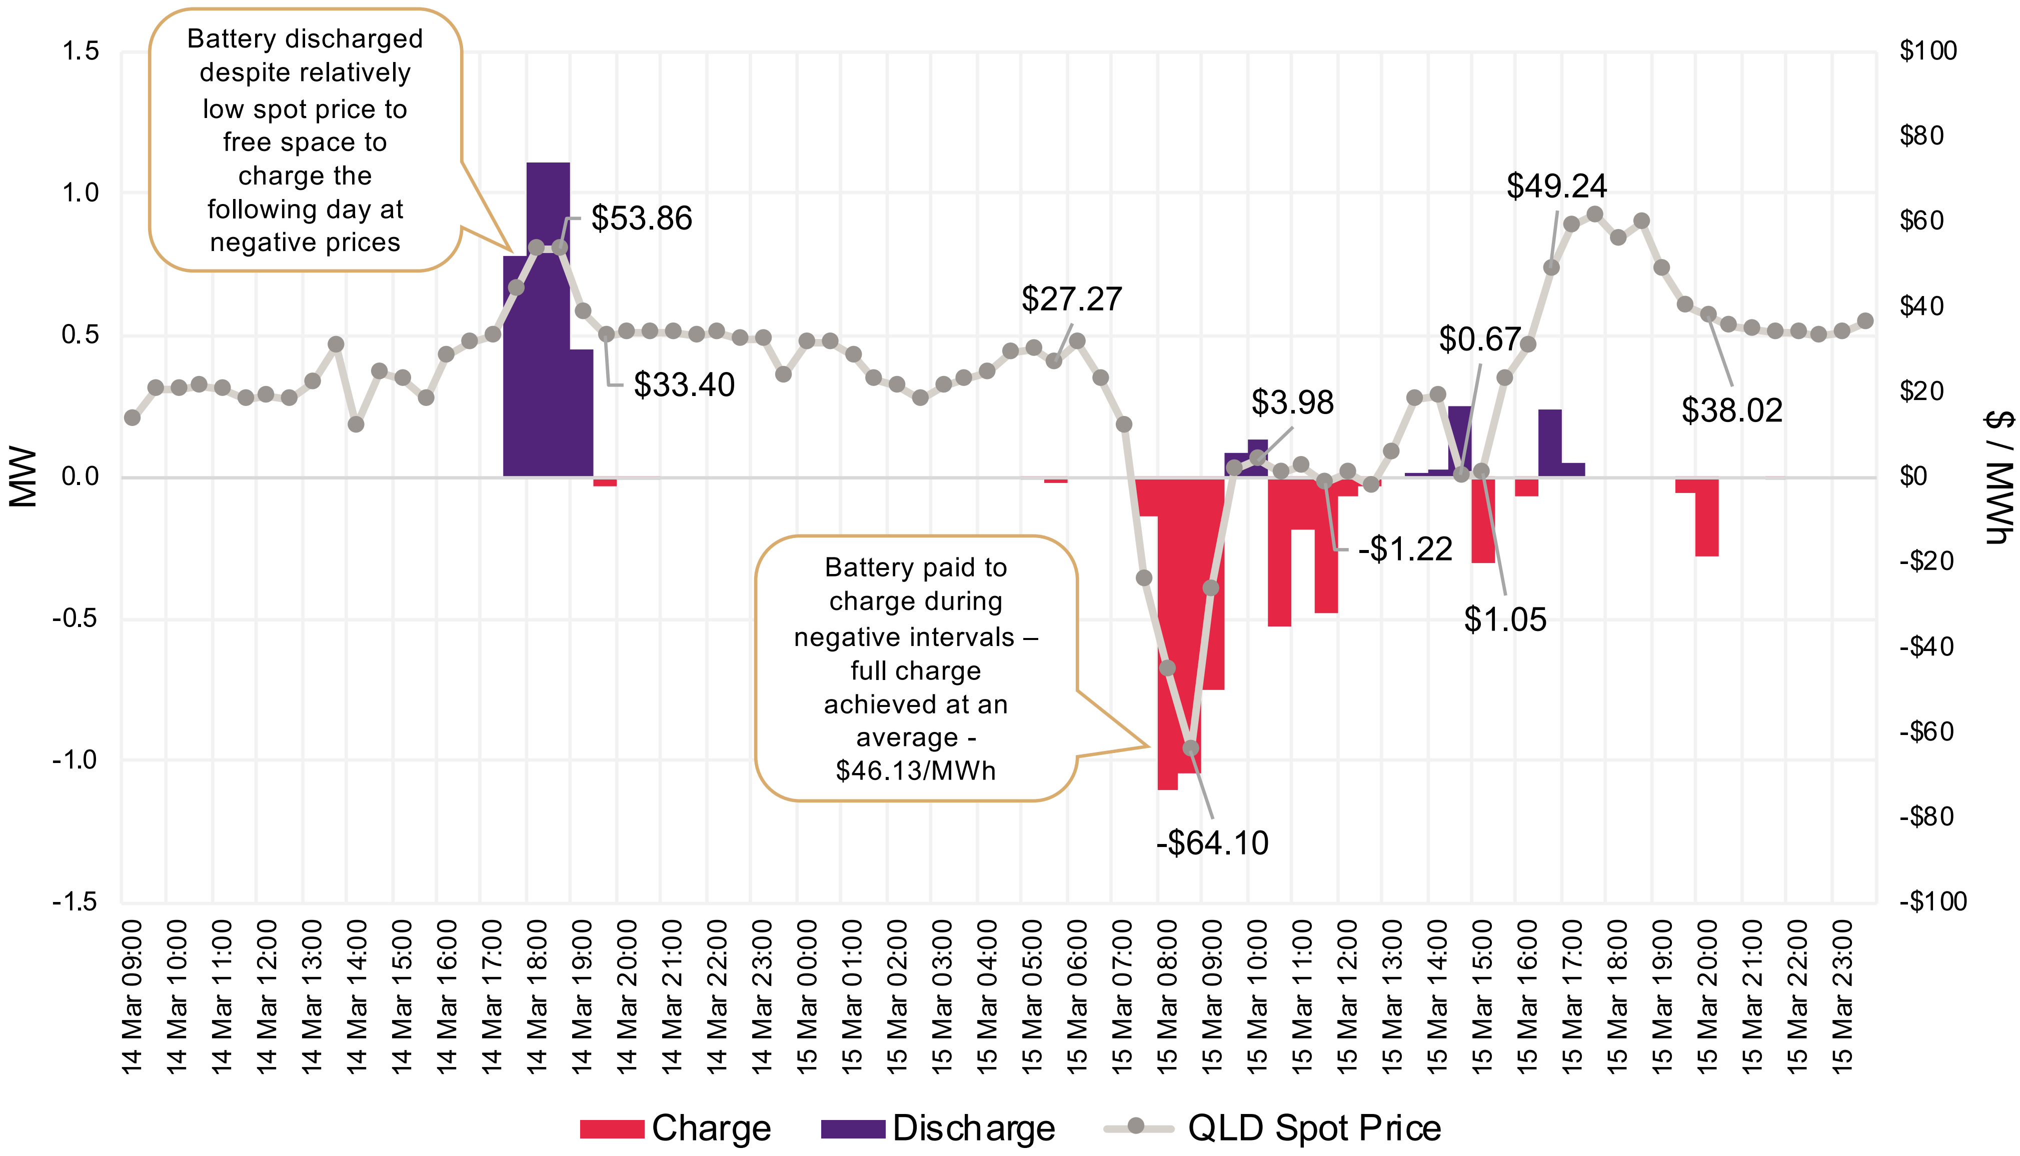

Calculation Methodology

Arbitrage performance is calculated based on the battery's electrical metering data to determine the net volume of energy charged or discharged in each interval. This is then multiplied by each interval's published spot price. The configuration of the battery's metering is such that all auxiliary load (including the cooling system) is included within the values measured for energy consumed. This enables the round-trip efficiency of the battery to be calculated by dividing the total energy sent out by the total energy consumed. It is important to note that this calculation is influenced to some degree by the state of charge that the battery starts and ends each period with. For example, if the battery started the quarter empty but ended it full this is not a true 'apples with apples' calculation of round-trip efficiency, however, over a three-month period this phenomenon has a relatively minimal influence - less than 0.6% at most.

One important consideration when calculating the net value of energy charged and discharged is the ancillary energy charges incurred by operating the battery, some of which are unique to the fact that the battery is located behind-the-meter of a major load site. These costs form many of the components of the site's retail electricity bill beyond the simple wholesale cost of the energy used which is billed at the spot price. These ancillary energy charges are levied on a c/kWh basis and include items such as TUOS and DUOS consumption based charges, AEMO market fees, and LGC and STC charges. When charging the battery, UQ incurs additional ancillary energy charges than would otherwise be the case due to the volume of energy measured by the front door meter being increased. When the battery discharges, the volume of energy at the front door meter decreases, effectively 'reimbursing' UQ for the extra ancillary energy charges that were incurred during period of charging. However, these values do not balance out to zero due to the round-trip efficiency losses of the battery. The net cost of these ancillary energy charges per month are provided in Table 3.1 in the following section. Across the quarter, these worked out to an effective operating cost of $4.30/MWh – a figure that is required to be factored into the calculations of the minimum viable spread threshold. Note that the impact of ancillary energy charges has been accounted for in all arbitrage revenue figures presented in this report, but it has not been included in the figures presented for average charge price, discharge price etc. as these purely reflect the volume weighted spot price during charging and discharging.

It is important to note that the method used to calculated charge cost, discharge income, and net arbitrage revenue is inclusive of all energy consumed or generated by the battery regardless of the intention at the time. This means, for example, that energy generated by the battery discharging during an FCAS event is accounted for in the calculations undertaken for arbitrage revenue even though this was not the reason why the battery was discharging in that interval. As a result, figures presented for FCAS and virtual cap contract revenue do not account for the value of the energy used to charge or discharge the battery for these purposes as this would be double counting.

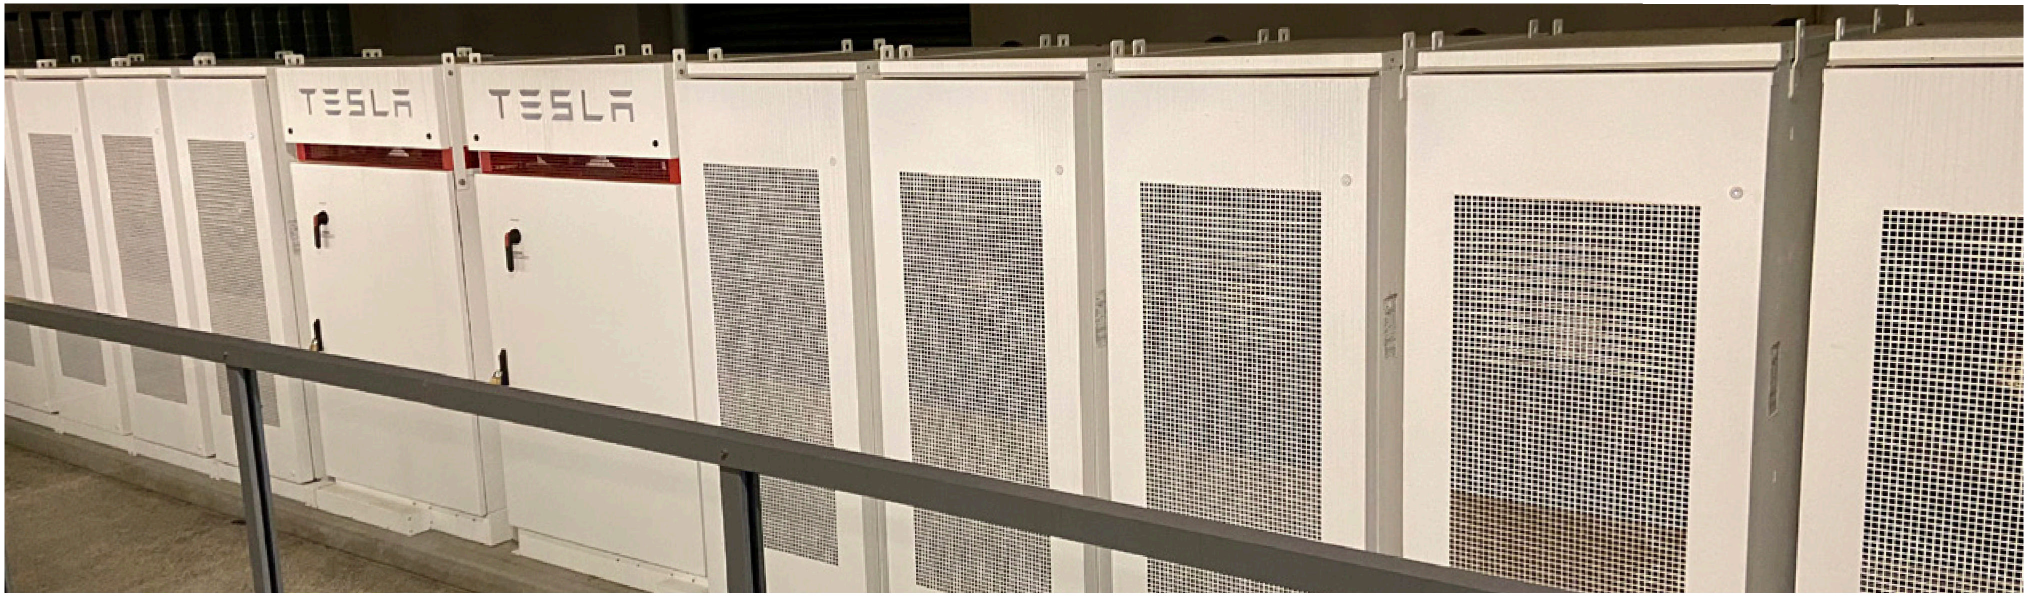

Table 3.1: Summary of arbitrage key performance figures by month

	January	February	March	Total
MWh charged	-35.42	-32.05	-40.23	**-107.70**
MWh discharged	31.09	26.95	32.94	**90.98**
Battery availability/uptime	96.7%	86.3%	99.5%	**94.3%**
Gross charge cost	-$2,108	-$1,423	-$1,122	**-$4,653**
Net ancillary energy charges	-$120	-$130	-$218	**-$469**
Gross discharge income	$7,776	$2,738	$3,131	**$13,645**
Total net revenue	$5,548	$1,185	$1,791	**$8,523**
Avg. charge price ($/MWh)	$59.51	$44.40	$27.89	**$43.20**
Avg. discharge price ($/MWh)	$250.11	$101.60	$95.05	**$149.98**
Avg. spread ($/MWh)	$190.60	$57.20	$67.16	**$106.77**
Time weighted avg. QLD spot price	$66.79	$53.81	$41.27	**$53.96**
Charge price % below QLD spot price	-10.9%	-17.5%	-32.4%	**-19.9%**
Discharge price % above QLD spot price	274.5%	88.8%	130.3%	**177.9%**
Charge price fraction of QLD spot price	0.89	0.83	0.68	**0.80**
Discharge price fraction of QLD spot price	3.74	1.89	2.30	**2.78**
Capacity factor (discharge only)	3.8%	3.5%	4.0%	**3.8%**
Capacity factor (charge + discharge)	8.1%	7.7%	8.9%	**8.3%**

Q1 2020 Performance

Full results and key performance figures broken down by month are provided in Table 3.1, noting that charge energy and costs are represented as negative values. Financial results for the arbitrage revenue stream are presented in Figure 3.4 which displays the battery's charge cost (inc. ancillary energy charges), discharge income, and net revenue by month.

The battery earned $8,523 in net revenue from arbitrage across Q1. This underperformed business case expectations by around 2%. This was driven by a range of factors but was predominantly related to total discharge volume being well below that forecast in the business case. A primary contributor to this was a physical fault with one out of ten battery packs (further discussed in section 4.1) that meant available storage capacity was reduced for the full duration of the quarter. Accounting for the 0.185 MWh that is required to remain reserved at all times for FCAS response (10 minutes at full power), this left only 1.75 MWh of capacity available for arbitrage (and virtual cap) purposes throughout Q1 2020 – 10% less than what had been assumed in the revenue forecasts which were based on all ten packs being available.

The second factor leading to this volumetric underperformance was higher than forecast outage durations. Revenue forecasts assumed a battery uptime/availability of 98% - equating to around 43 hours of total outage time. In reality however, availability only reached 94.3% across the quarter. This resulted in around 124 hours (or 5 cumulative days) of outage, and was largely driven by nuisance tripping issues during January and February as further discussed in section 4.1.

In total, the battery discharged an average of 1.00 MWh per day across the quarter – noting that volume discharged is the most relevant figure for the purposes of this analysis. Based on full battery pack availability, one full discharge per day with 98% availability (as forecast) should have resulted in an average discharge volume across the quarter of around 1.92 MWh per day. This represents a total volumetric underperformance of 48% in Q1.

It is important to acknowledge that volumetric outcomes were also to some degree driven by the underlying dynamics of the market. This includes the availability of minimum spreads each day (particularly towards the back end of the quarter) as well as DRE's effectiveness at being able to predict and achieve these in the context of the challenges discussed in section 4.3.

Figure 3.4: Charge cost, discharge income & net revenue by month

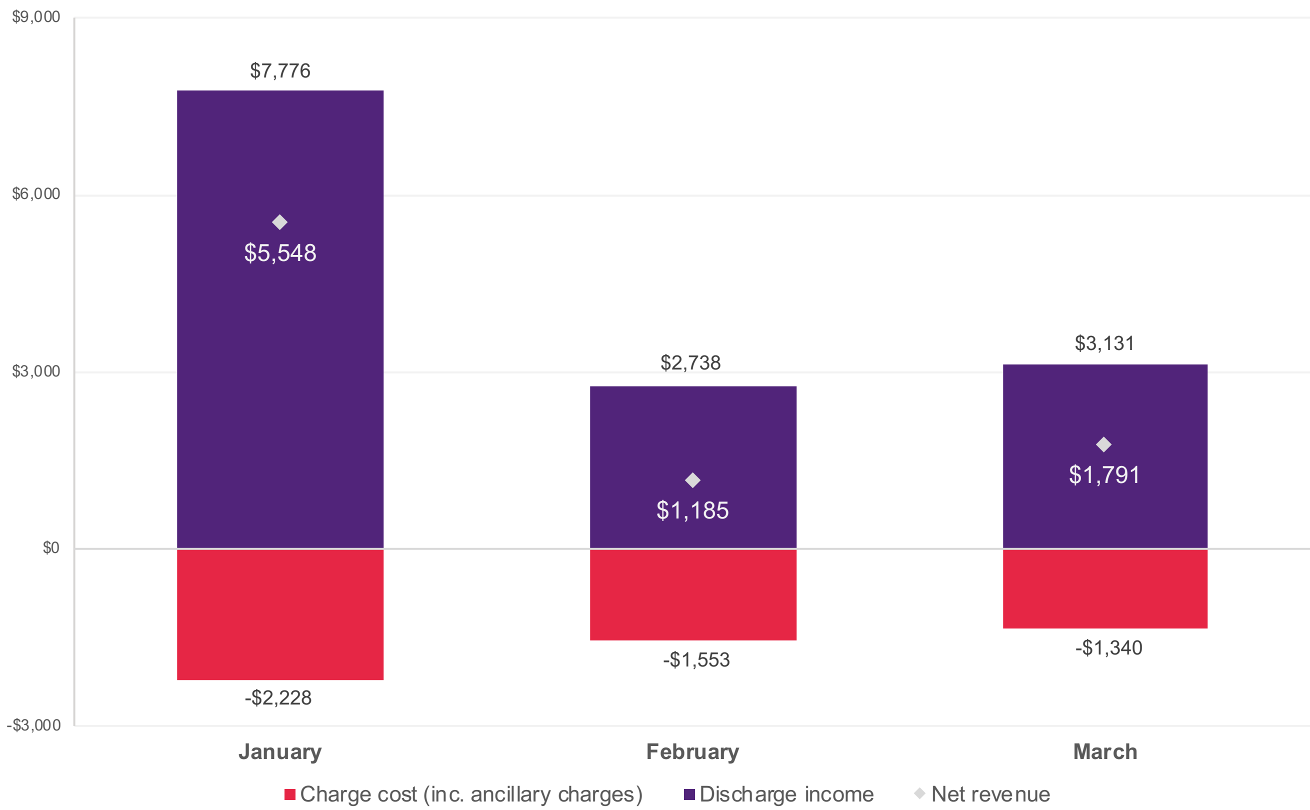

The issues with lower than forecast volumes being discharged were to a large degree offset by substantially higher than forecast spreads being achieved for the volumes that were dispatched. This resulted in the battery only marginally missing the quarter's target revenue figure. If the same average spread would have been achievable for full dispatch volumes (one cycle a day, ten packs online, 98% availability), arbitrage revenue for the quarter would have been around $18,600 - a 114% overperformance of business case forecasts. While this is an overly optimistic extrapolation, it nonetheless highlights the importance of addressing some of the underlying issues that were identified in Q1 that caused dispatched volumes to be reduced.

The battery's Q1 arbitrage revenue was heavily skewed towards January in which net revenue was almost double the value earned in February and March combined. This is a reflection of the underlying pricing volatility in the NEM during January compared to the other months. This is most clearly represented by Figure 3.5 which shows that although cumulative arbitrage revenue gain was steady across the full quarter, large 'jumps' in this revenue occurred as a result of high-priced intervals on a handful of days in January.

Figure 3.5: Cumulative arbitrage net revenue

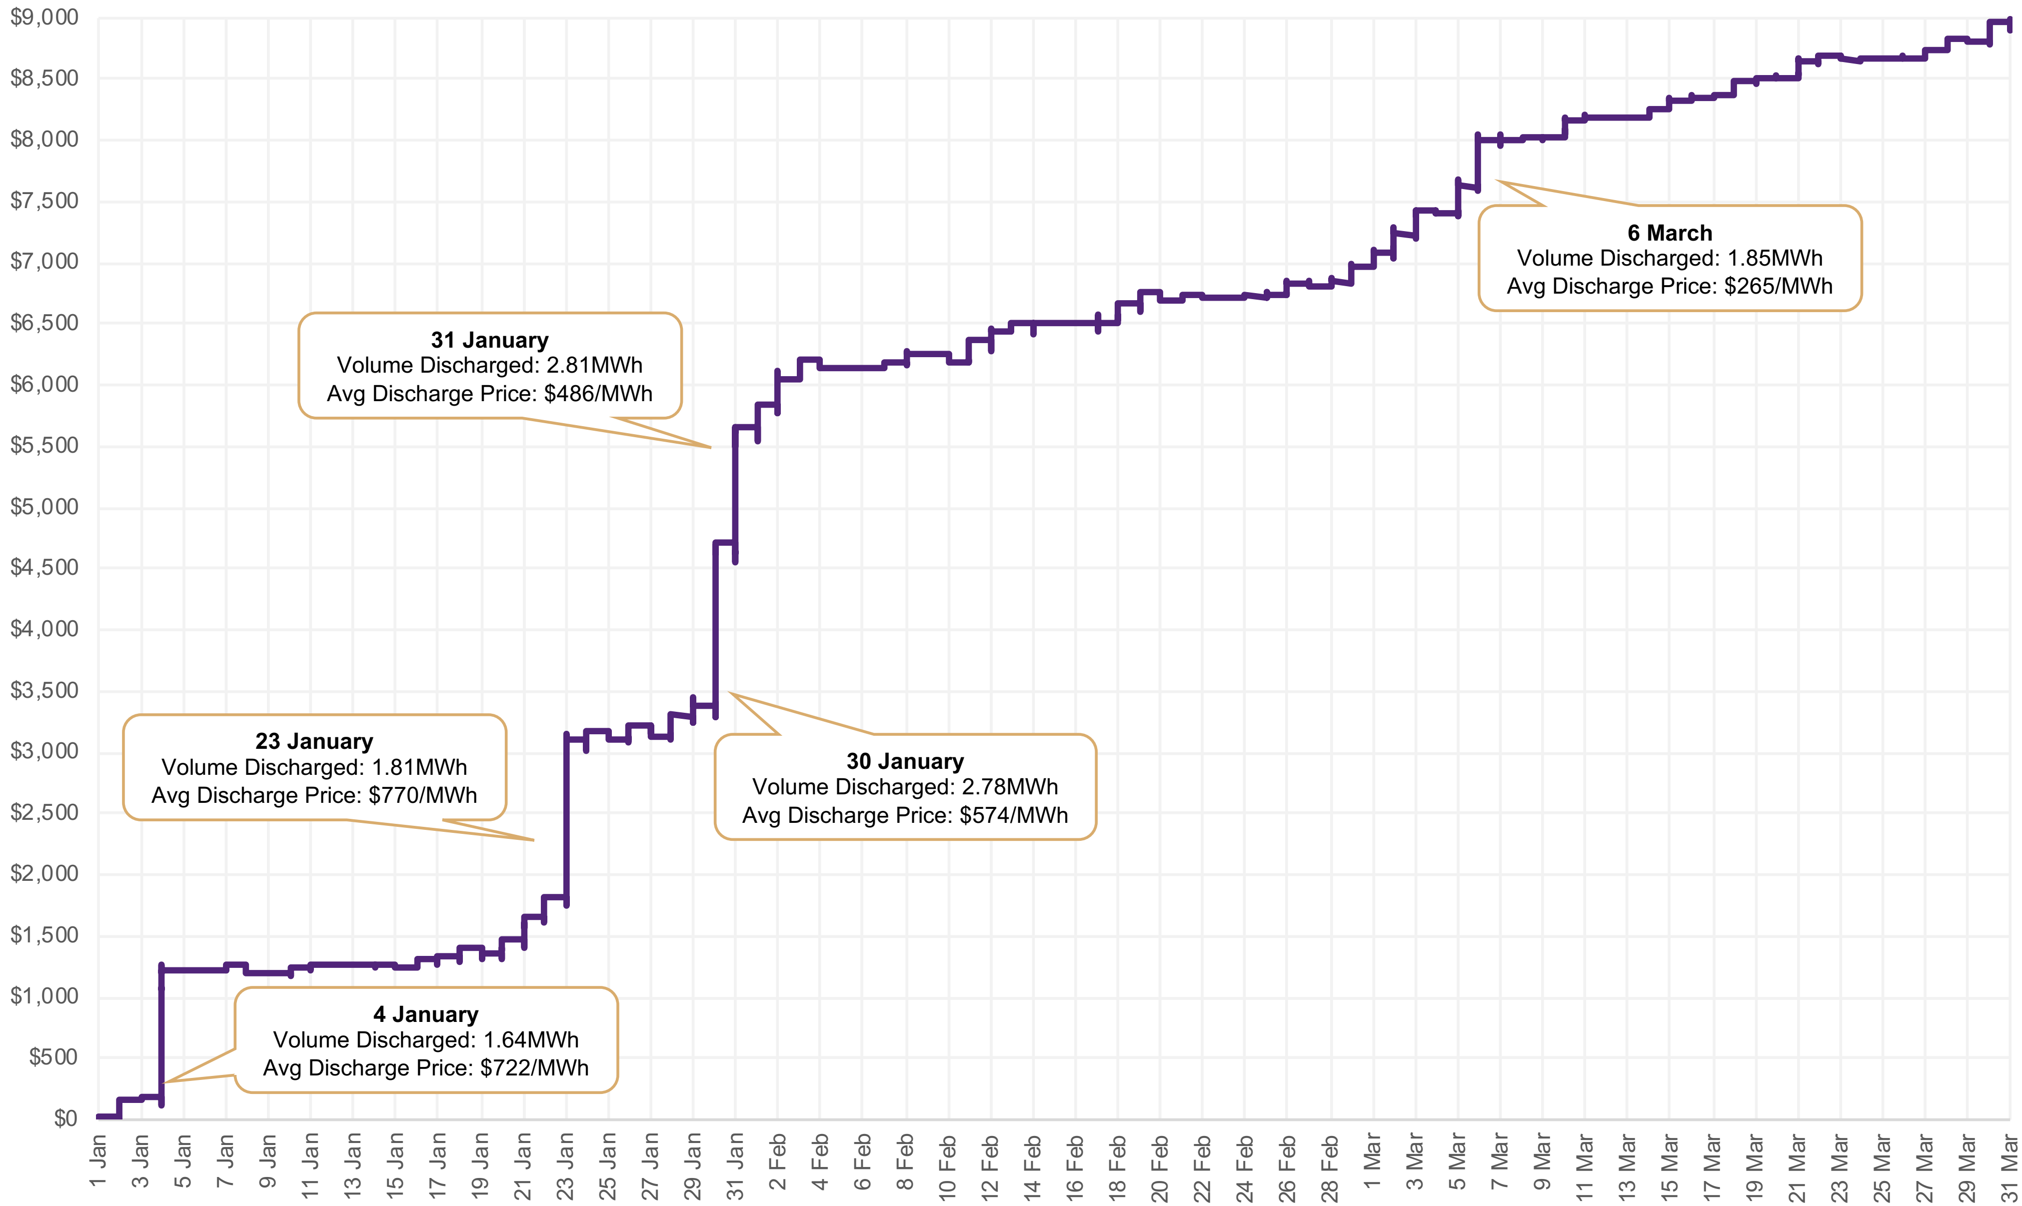

3.3 Frequency Control Ancillary Services (FCAS)

One of the key characteristics typical of batteries versus other forms of energy storage is their incredibly fast response times. The UQ battery has been measured during commissioning tests to be able to respond from zero to full discharge as fast as 200 milliseconds, or to swap from full charge to full discharge in around 400 milliseconds. This makes it an ideal technology for participation in the NEM's various FCAS markets.

As a behind-the-meter asset less than 5 MW in size, the UQ battery currently participates in only three contingency FCAS markets – Raise 6 seconds, Raise 60 seconds, and Raise 5 minutes. These markets require capacity to be available to respond to a drop in system frequency below the dead band threshold of 49.85hz within the stated timeframes, and for this response to be sustained until frequency is restored to the normal operating range (usually within 10 minutes). With a sub one second response time being achievable, the battery is able to participate in any of the three markets and may have different volumes of capacity bid across each different market, depending on the desired bidding strategy. It is important to note that FCAS providers are paid for every interval in which they are available to respond to a frequency event, and that revenue is not linked to the number of times or the duration of such events that occur.

To participate in the contingency FCAS market, UQ has partnered with Enel X (formerly Enernoc). UQ and Enel X's partnership extends back to February 2018, when UQ's 600 kW pilot lithium ion battery at the Gatton campus was the first behind-the-meter battery in Australia to participate in the FCAS market, joining the grid only two months after the Hornsdale Power Reserve commenced operation. This relationship has continued with the St Lucia battery, whereby Enel X is responsible for bidding an overall portfolio of contingency FCAS raise capacity into the market, of which UQ's battery is a key component.

An example of the battery's contingency FCAS performance is illustrated in Figure 3.6. This shows (at a four second time scale) events that occurred on the afternoon of Friday 6 March, when at 4:26pm (QLD time), Unit 4 of the Loy Yang A brown coal power station in Victoria suddenly tripped offline, resulting in the sudden loss of 558 MW of generation and causing the mainland NEM frequency to dip as low as 49.751 Hz in the 4 second data captured by UQ – well below the contingency FCAS trigger point of 49.85 Hz. The battery sensed this frequency deviation immediately and was able to switch from charging at a rate of 919 kW to discharging at a sustained rate of 1,099 kW - providing a total of 2.02 MW of 'generation' into the network to help arrest the fall in frequency. This response was sustained for 304 seconds, after which frequency restored to within the normal operating range.

Figure 3.6: Contingency FCAS performance on the afternoon of Friday, 6 March following Loy Yang A unit trip

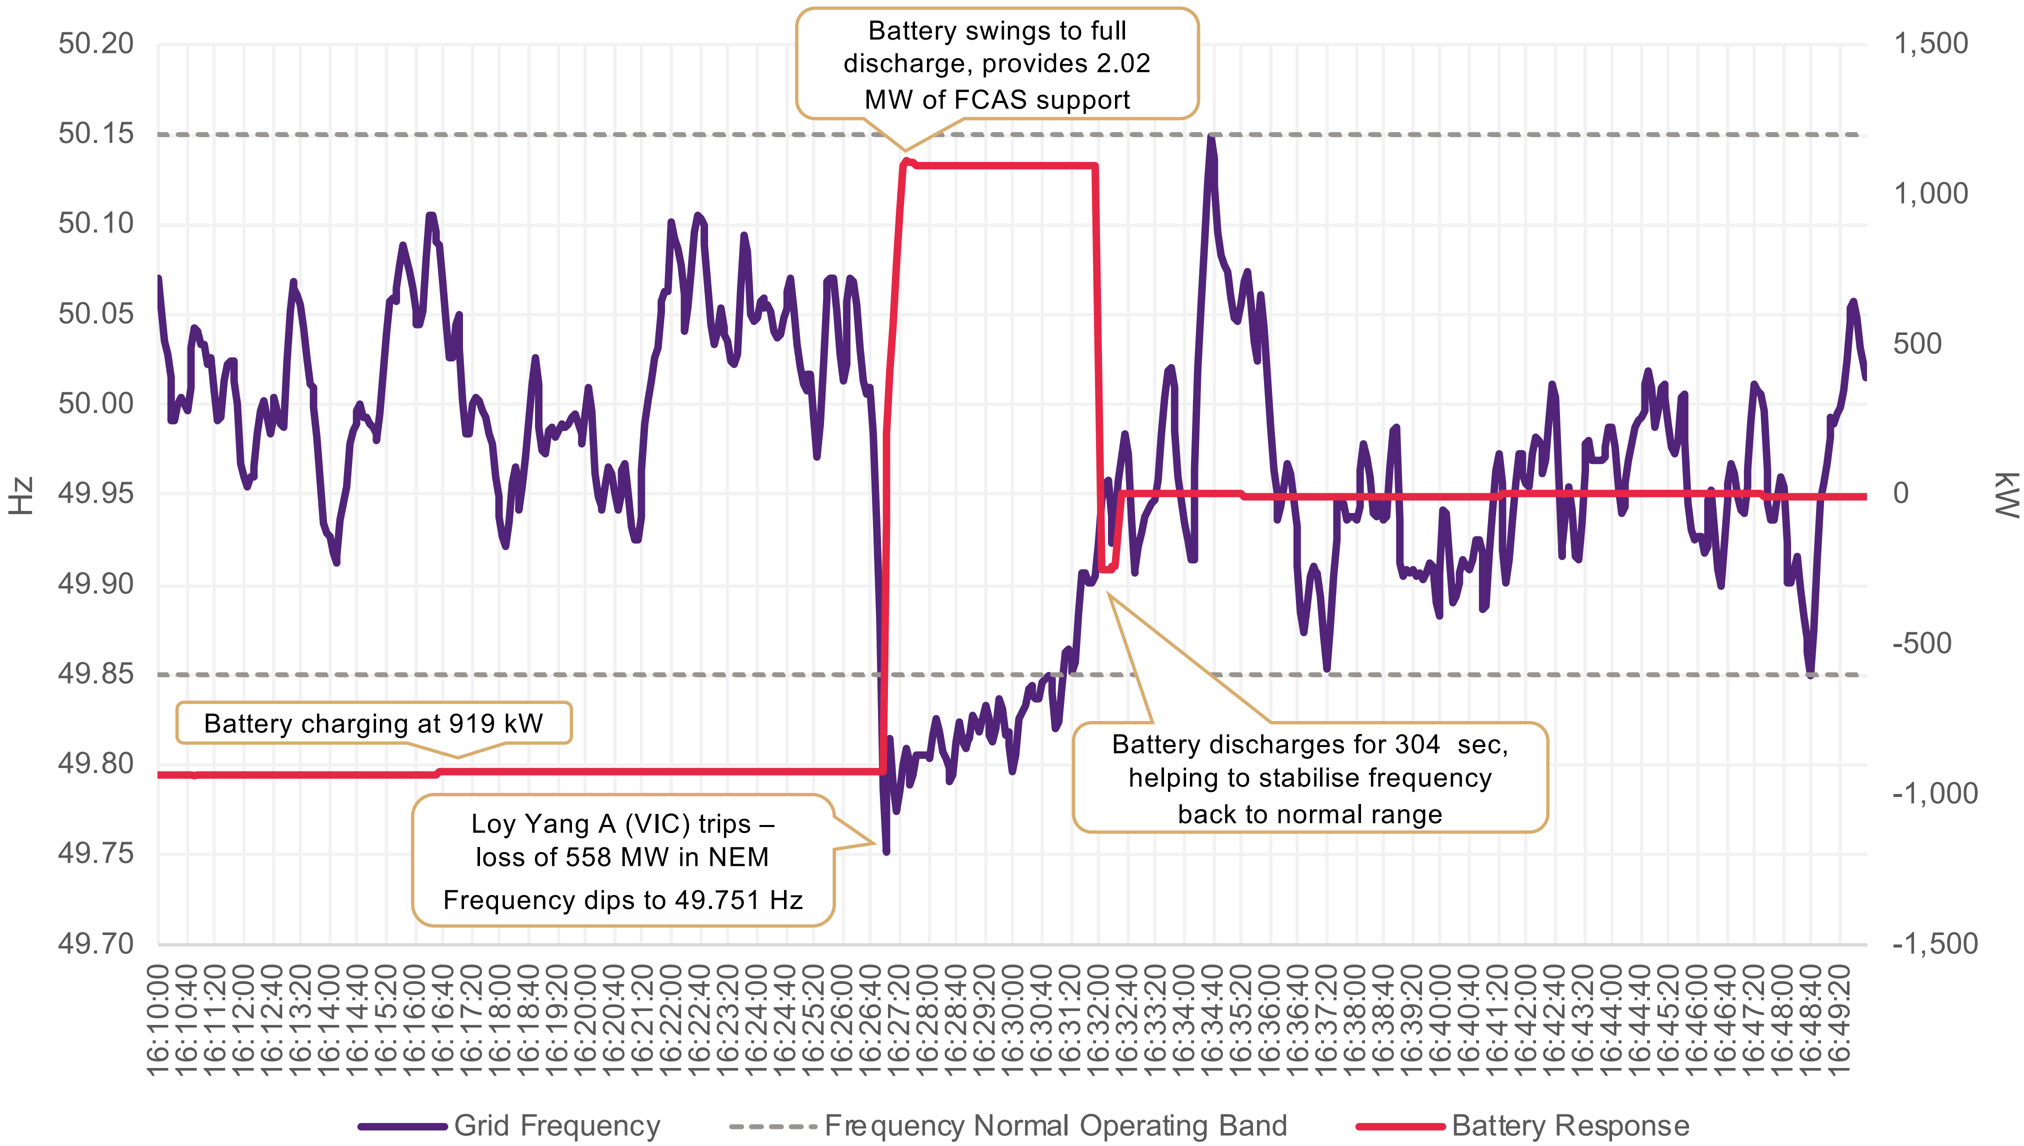

Calculation Methodology

The revenue earned by the battery for FCAS services is a function of the volume of capacity bid into the FCAS markets, the way this volume is bid across each of the three contingency raise markets, the prevailing market prices, and a commercial-in-confidence revenue sharing agreement between UQ and Enel X. Revenue figures for each month were provided by Enel X and verified with UQ's metering and the battery's SCADA data.

As discussed in section 3.2, no accounting for the cost or income of the energy component of battery charging and discharging for FCAS purposes is made here, with these figures already being captured in the values provided for the arbitrage revenue stream.

Q1 2020 Performance

In total, the battery earnt $46,000 in FCAS revenue across Q1 2020. It responded to a total of 12 frequency events, with the cumulative duration of response across the quarter totalling 42 minutes and 16 seconds. Full results broken down by month are provided in Table 3.2, noting that revenue figures have been rounded due to commercial considerations.

As seen in Table 3.2, FCAS revenue in January was almost twice that of February and March, as were the number of FCAS events and total response time. This was largely driven by a combination of natural disasters that struck the NEM during the month. Firstly, the unprecedented Black Summer bushfires led to the electrical separation of the NSW and VIC regions (and Snowy sub-region) on the afternoon of 4 January as a result of damage to transmission infrastructure. This resulted in a sustained period of FCAS (and energy) pricing close to the market ceiling as the supply and demand balance across the NEM was disrupted.*

Then in the early afternoon of 31 January, storm damage to transmission towers led to South Australia 'islanding' from the rest of the NEM, again causing an extended period of price volatility in the FCAS markets.** As an indication of the effect that these events had on FCAS revenue, average pricing for the Raise 6 second market during January was over four times higher than in March, and nearly three times higher than in February. This helped contribute to overall FCAS revenue in Q1 2020 exceeding business case assumptions by over 50%.

Table 3.2: Summary of FCAS key performance figures by month

	January	February	March	Total
Total revenue	$23,000	$13,000	$10,000	**$46,000**
Total # of FCAS events	6	3	3	**12**
Total duration of events (seconds)	1,632	128	796	**2,556**
Avg event duration (seconds)	272	43	265	**213**
FCAS MWh discharged	0.083	0.013	0.081	**0.177**

*__Read more:__ http://www.wattclarity.com.au/articles/2020/01/bushfires-trip-victonsw/

**__Read more:__ http://www.wattclarity.com.au/articles/2020/02/31jan2020-howdidthelightsstayoninsa/

3.4 Virtual Cap Contract

As a participant in the wholesale electricity spot market, UQ is required to develop risk management strategies to the potential impacts of market volatility. One option available to help manage this risk is the use of 'cap' contracts. These financial products act as a form of insurance against extreme market prices. The buyer of the cap contract pays a 'premium' (typically expressed as $/MWh for a fixed volume) and is 'paid out' if and when market prices exceed a set threshold. This threshold is typically $300/MWh, and the payout is calculated on the difference between this level and the spot price for the relevant interval. For example, if the spot price in a half hour interval was $2,500/MWh, the cap contract holder would be paid out $2,500 subtract $300 (i.e. $2,200/MWh) multiplied by the volume of the cap contract held. At the end of a defined period (e.g. a quarter), the cost of the premium minus the revenue from any intervals where the contract paid out is the net value of the cap contract to the holder.

As a behind-the-meter asset that is able to respond quickly to market price spikes, the UQ battery is able to replicate the risk management function performed by a financial cap contract at least partially. This 'virtual' cap contract function works by the battery discharging stored energy during intervals where prices spike beyond a set threshold (e.g. $300/MWh). This then reduces UQ's load by 1.11 MW and thus UQ's exposure to the high market price by the same volume.

The primary shortfall of a virtual cap versus a financial cap is that it is unlikely the battery will be able to respond to every price spike in a way that provides coverage for the full half hour interval. This is largely a function of the NEM's current 5/30 settlement dynamics, where dispatch spot prices are set every 5 minutes, but trading prices (which are used for financial settlement) are based on the half hourly average of six dispatch prices. This can result in scenarios where the 5-minute dispatch price unexpectedly spikes mid-way through a 30-minute interval, leaving UQ partially exposed to the high price, even if the battery responds immediately and discharges for the remainder of the half hour. The virtual cap contract is also subject to many of the same constraints as the battery's other services – namely physical limitations on the duration of energy storage. This means that coverage may fall short during periods of prolonged market volatility as the battery's stored energy depletes. On the other hand, the primary benefit of a virtual cap is that once the capital cost of the battery has been expended, there is no ongoing premium payable for this service like there would be for a financial cap.

Calculation Methodology

The net value of the virtual cap contract to UQ needs to be calculated with reference to the appropriate avoided cost of using this in lieu of a financial cap. This requires the net value of a hypothetical financial cap contract to be calculated for comparison purposes. The first step to do this is to determine the gross cost of its premium over the quarter. In this case, a rate of $12.37/MWh has been used, being the average price of Q1 2020 cap contracts from the ASX Energy exchange as traded during the three months preceding the start of Q1 2020. Gross income from the pay out of the contract in all intervals greater than $300/MWh is then subtracted from this premium to determine the net value of the financial cap over the quarter.

The value of a virtual cap can then be determined by comparing it to the financial alternative over the same period. This is done by first calculating the costs incurred by the virtual cap in the periods where full 'cover' was not provided. For example, in an interval where the spot price was $1,500/MWh, a financial cap would have paid out $1,200/MWh multiplied by the cap volume (in this case, 1.11 MW or 0.555 MWh per half hour). If in the same interval the battery was only able to respond and discharge for 20 out of the 30 minutes, there would be a 0.185 MWh shortfall which would carry with it a cost 'penalty' of $222 for the interval.

The overall value from the battery's virtual cap service is then calculated from the difference between the net cost of the hypothetical financial cap of the same size subtract the sum of the virtual cap cost penalties over the same period. Refer below for the specifics of the calculation for Q1 2020.

Note that no cost or revenue from the energy associated with charging or discharging the battery is attributed to the revenue figures provided for the virtual cap contract service. This is due to these values already being captured in the methodology which calculates the total revenue of the arbitrage value stream. As a result, including the value of energy brought or sold due to the virtual cap contract service would be double counting.

Q1 2020 Performance

It has been calculated that a 1.11 MW financial cap contract would have represented a net cost to UQ of $22,804 over the quarter. This is comprised of $29,995 of gross premium cost subtract $7,190 of income from the payout of the contract across the 18 intervals that exceeded $300/MWh during the quarter.

During this same period, the virtual cap resulted in a cost to UQ of $3,389 based on the shortfall in its coverage of high price intervals. In total, the virtual cap provided 5.79 MWh of coverage out of an overall total of 9.99 MWh of exposure across the 18 intervals – equivalent to coverage of 58% on a volume basis. An alternative performance metric is to assess the percentage of cover provided on a financial basis. That is, did the periods of missed coverage skew towards being relatively higher or lower priced? This is an important question, as faced with only limited discharge duration the battery should theoretically prioritise covering those intervals that are higher priced. In reality however - and for the reasons discussed further below - coverage on a financial basis was lower than that on a volume basis, at only 53%.

Overall the virtual cap service delivered **$19,415** in net value to UQ across Q1 2020 compared to a financial cap alternative – the $22,804 avoided cost of a financial cap subtract the $3,389 penalty of missed coverage. Full results broken down by month are provided in Table 3.3, noting that costs are represented as a negative value.

As seen in Table 3.3, market volatility as measured by the incidence of prices >$300/MWh was largely confined to the start of the quarter, with 15 out of 18 total intervals occurring during January. Of note, this resulted in the net value of the virtual cap being marginally negative as of 31 January. This was due to the income vs. premium cost to date of the financial cap being relatively balanced at that point, as well as the cost of the battery's missed coverage during several high priced intervals. This trend is clearly illustrated in Figure 3.7 which shows the cumulative net value of the virtual cap (as per the above methodology) across the quarter. This highlights that the value to UQ was accrued in February and March where there were limited intervals >$300/MWh but the cost of the financial cap premium continued to be incurred, eroding the benefit of earlier payouts.

Table 3.3: Summary of virtual cap key performance figures by month

	January	February	March	Total
Financial cap gross premium	-$10,218	-$9,559	-$10,218	**-$29,995**
Financial cap gross income	$6,923	$27	$241	**$7,191**
Financial cap net value	-$3,295	-$9,532	-$9,977	**-$22,804**
# of intervals above $300/MWh	15	1	2	**18**
Max. potential exposure MWh	8.325	0.555	1.11	**9.99**
Max. potential exposure $	-$6,922	-$27	-$241	**-$7,190**
MWh covered by battery	4.231	0.455	1.101	**5.787**
MWh left exposed	4.094	0.100	0.009	**4.203**
% volume cover by battery	50.8%	82.0%	99.2%	**57.9%**
Cost of virtual cap missed coverage	-$3,381	-$5	-$3	**-$3,389**
% financial cover by battery	51.2%	81.5%	98.8%	**52.9%**
Net value of virtual cap vs. financial cap	-$86	$9,527	$9,974	**$19,415**

Figure 3.7: Cumulative net value of the virtual cap service across Q1

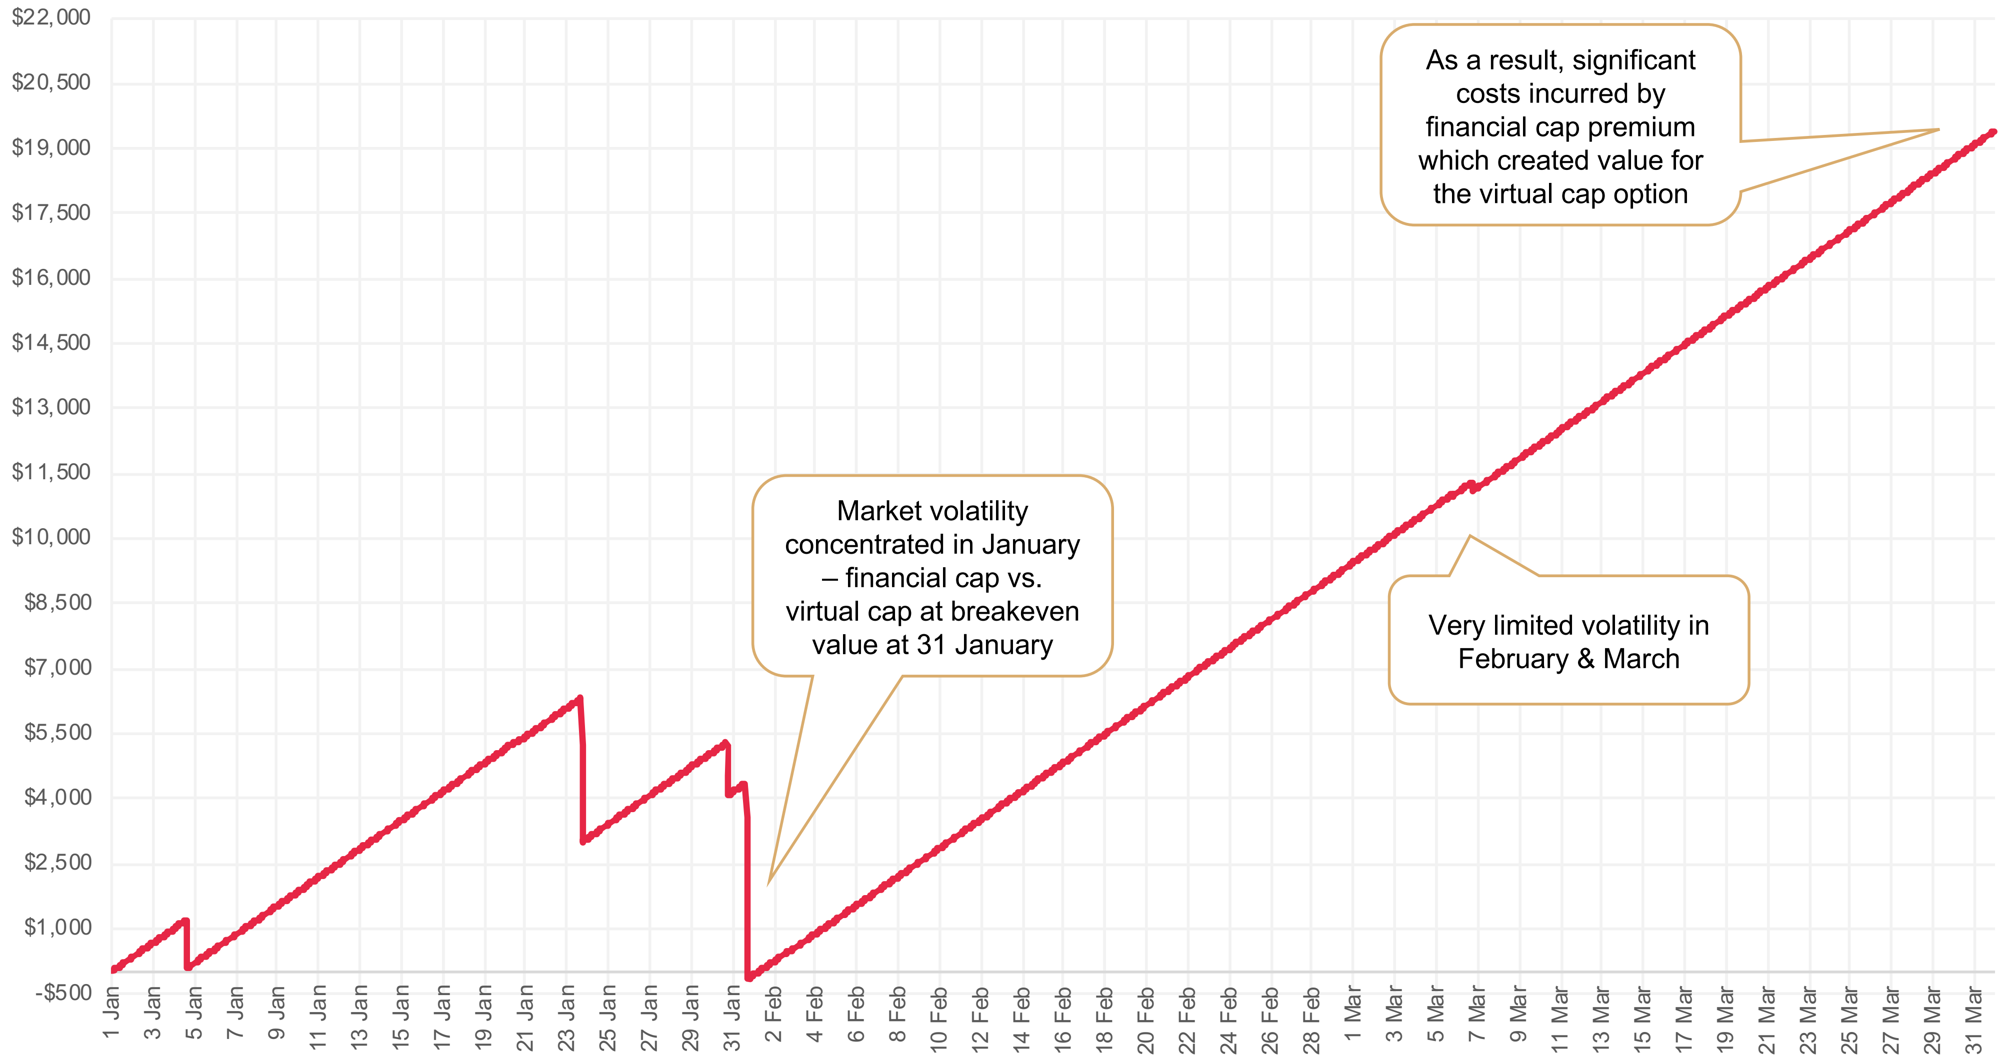

As shown in Figure 3.8, while overall volume coverage of 58% was achieved across the quarter, performance during individual intervals was highly variable. Most notably, the volume left uncovered during intervals on 23 January and 31 January exceeded 100%. This is due to the battery charging for at least part of the interval where the overall outcome should have instead been to discharge. Both of these events occurred as a result of erratic forecasts and a low state of charge that drove that battery to charge at what was at the time seen as a 'cheap' price in order to be ready for higher priced intervals that were forecast in the near future but did not eventuate. These outcomes provide a good illustration of the challenges faced by DRE's control algorithm in optimising the battery's performance during volatile periods and which are further discussed in the full report. As these two instances occurred when market pricing was $2,296/MWh and $1,610/MWh respectively, they had an outsized influence on the overall percentage of financial cover provided by the battery across the quarter. Excluding these two intervals the battery would have provided 66% volume cover and 73% financial cover.

Figure 3.8: Virtual cap volume cover performance during >$300/MWh intervals

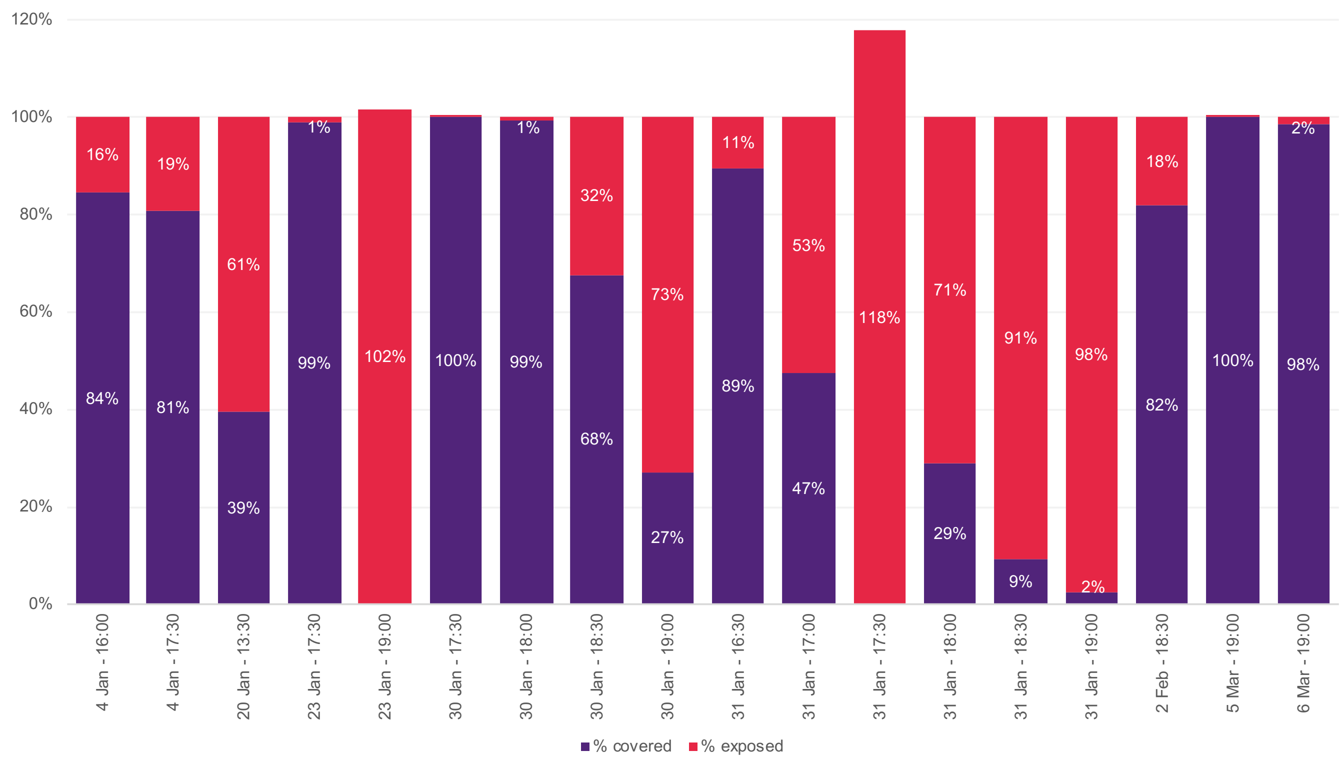

3.5 Peak Demand Lopping

As a large energy consumer, UQ's monthly electricity bill for the St Lucia campus contains a mechanism to charge for the site's peak demand – measured as the highest average kVA reading over a half hour interval. The business case for the battery included the value that may be able to be derived from using it to strategically 'lop' the top off each month's peak demand level. It was recognised that UQ's typical peak demand each month is often prolonged (i.e. not a sudden 'spike') which requires a careful trade-off between targeted kVA reduction and the duration of energy storage available. For example, it is unrealistic that UQ could reduce monthly peak demand by 1,110 kVA, as the duration required to make a difference between the peak and the next highest demand intervals is likely to be longer than the 2 hours of sustained discharge that would be available at this power set point. Instead, peak demand lopping of around 500 kVA over a cumulative total of 4 to 4.5 hours in a day is more likely be required.

The control strategy for this functionality was not finalised in time to enable its deployment during Q1 2020. In order to finalise this work, the ability to forecast campus peak demand utilising weather forecasts from the Bureau of Meteorology needs to be implemented. This has previously been trialled with a high degree of success and is now at the stage of being refined ready for deployment. Most significantly however, the ability to forecast and 'lop' peak demand needs to be successfully integrated with the control algorithm for arbitrage in order to co-optimise the two strategies which may otherwise sometimes cause the battery to take contradictory actions. This is discussed further in section 5.1.

Implementation of the site demand forecasting capability is also essential for ensuring that the battery's operation does not negatively impact campus peak demand. This was emphasised by an unfortunate example in Q1 2020 whereby DRE's arbitrage control algorithm directed the battery to charge ahead of forecast high spot pricing but during a period whereby the site was seeing its highest monthly demand at the same time. In total, this resulted in campus peak demand for the month being 854 kVa above where it would have otherwise been, as shown in Figure 3.9. This resulted in a cost penalty of $2,357.

As the cause of this issue was known and a solution to address it exists, this cost penalty has not been deducted from the battery's overall Q1 revenue figures.

It is worth noting that the demand charges paid by UQ at the St Lucia campus are significantly lower than those paid by many other commercial energy users. This is driven by the size and voltage of the connection, as the St Lucia campus has its own embedded distribution network beyond the front door 11kV connection point. For example, the variable monthly peak demand charge at St Lucia for the 2019/20 financial year is $2.76 per kVA per month*. This compares to rates for smaller sites at $15.80 per kVA per month such as under the Energex 8100 tariff.

This relatively unique disparity of UQ's demand charges inherently results in a different business case for large scale storage behind-the-meter. At a site with more typical demand charges, the revenue expected from peak demand lopping could be more than five times higher, and likely a significantly higher portion of revenue than arbitrage. This also reflects the fact that at present, many commercial and industrial energy users do not have wholesale spot price exposure. This means that value UQ is able to derive from arbitrage and virtual cap contract services is unlikely to be as applicable to them. This highlights why UQ is committed to developing and trialling methods of peak demand forecasting and lopping, despite the value in our specific application likely being relatively minor.

*Excludes capacity charges that are set annually on a rachet basis and are worth an additional $2.77/kVA per month. Battery driven peak demand lopping would theoretically also result in a reduction of these charges, although further work is required to demonstrate the feasibility of this in practice.

Figure 3.9: Metered kVA demand (with battery) vs. actual underlying kVA demand (without battery) on 20 and 21 January

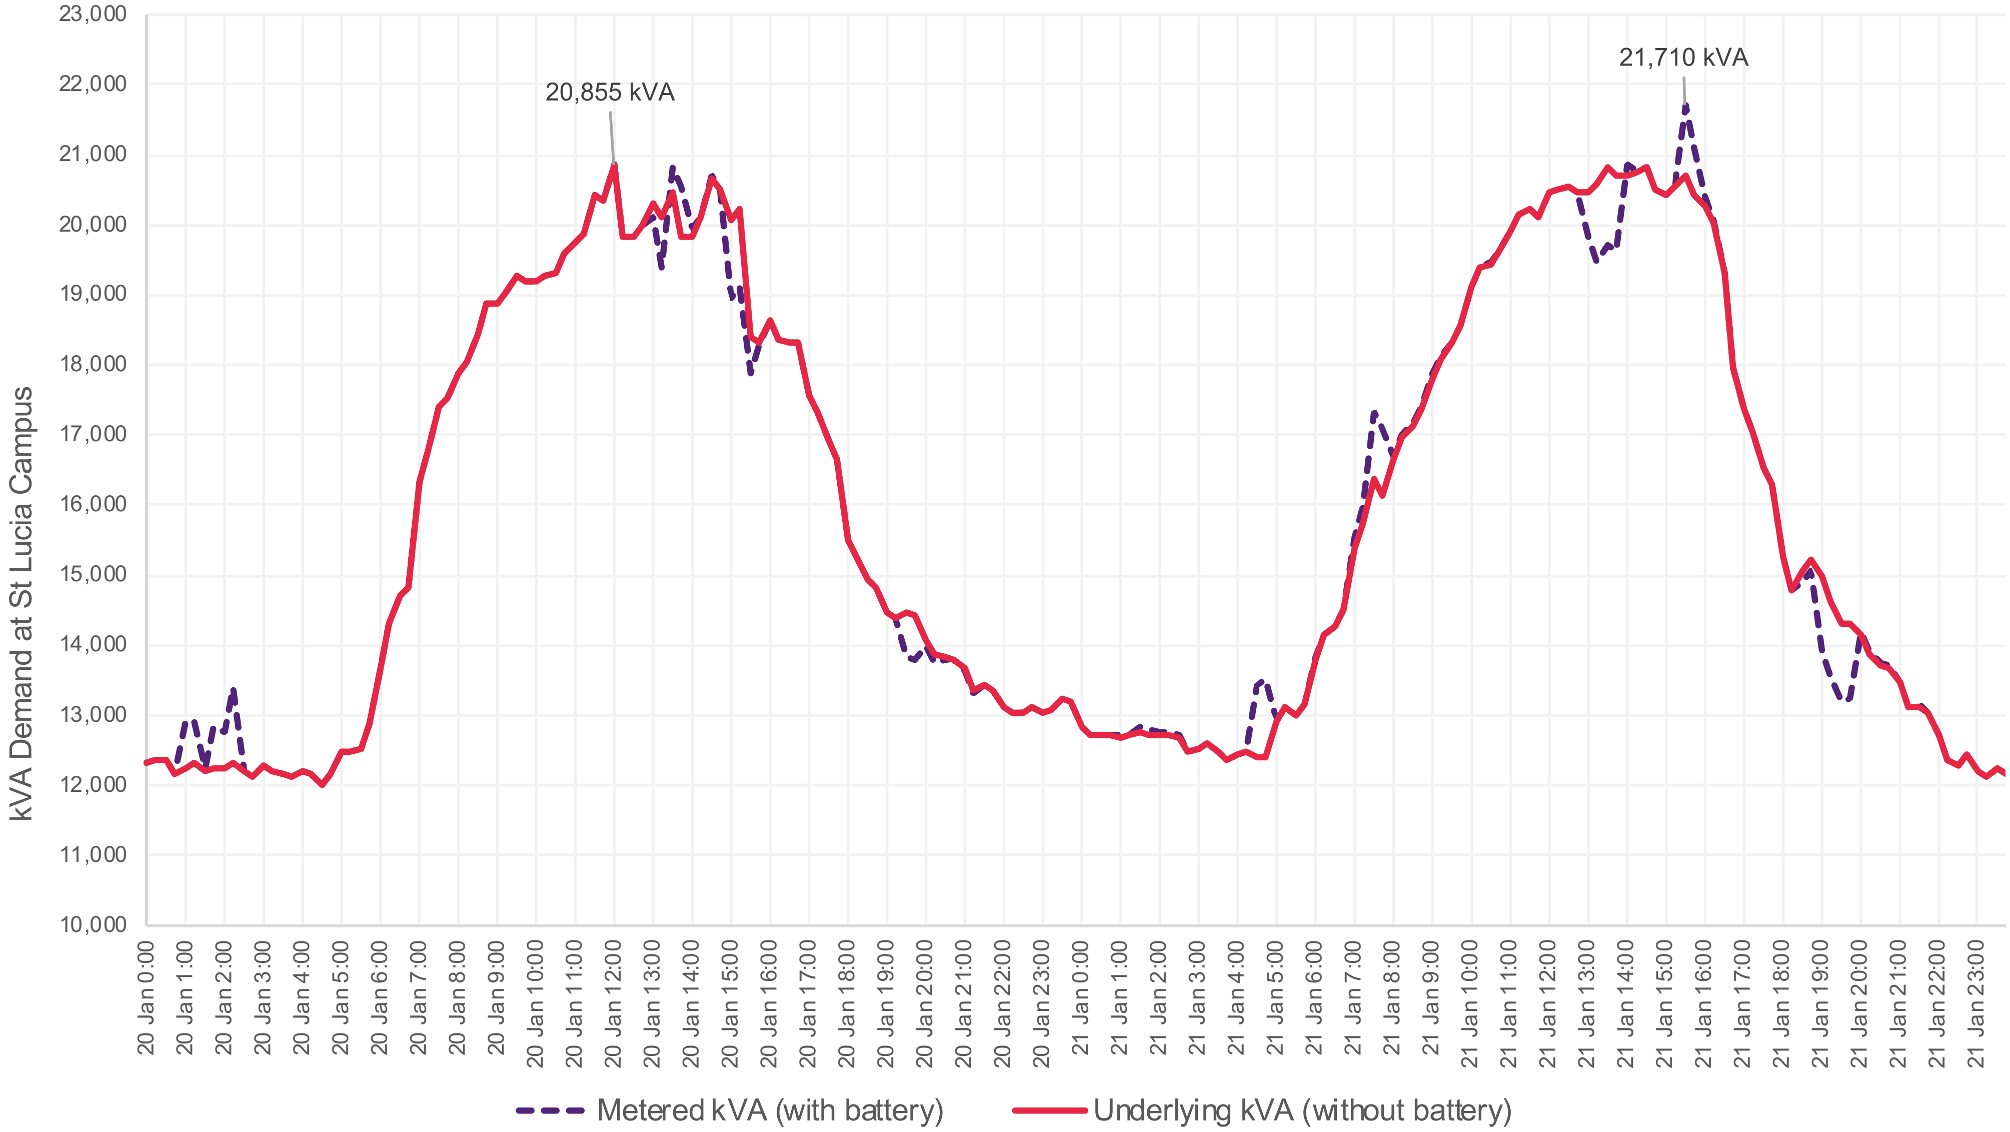

3.6 Comparison to Wivenhoe Pumped Hydro

As well as comparing Q1's results against business case forecasts, it is possible to assess the battery's performance against the major energy storage asset in Queensland – the Wivenhoe Pumped Hydro Power Station. With a pumping (charge) capacity of 2 x 252 MW and a generation (discharge) capacity of 2 x 285 MW, Wivenhoe is substantially larger than UQ's battery. As a pumped hydro plant, its storage duration characteristics are considerably different than a battery, however it does not share the battery's same ability to respond to market events in under one second. Most notably though, Wivenhoe is a major generation asset that is operated by a specialised energy company (CleanCo) with a dedicated trading team. This is as opposed to the UQ battery that is operated autonomously via DRE's control system. These factors accordingly make for an interesting comparison.

It is important to note that this comparison is not ideal, as while performance data is publicly available, it is not possible to know why Wivenhoe traded the way it did. For example, it is likely that the CleanCo trading team also needed to consider how to operate Wivenhoe as part of their overall portfolio position in relation to aspects such as forward contracts that had been sold. Capacity factor and nameplate revenue figures in the below analysis have assumed that pumping and generating at Wivenhoe occurs on a duty + standby arrangement, whereby only one pump or generator typically operates at one time.

Table 3.4: Comparison of Wivenhoe and UQ battery across key metrics

	Wivenhoe	UQ Battery
Charge MWh	50,219	107.70
Discharge MWh	33,433	90.98
Round trip efficiency	66.6%	84.5%
Gross charge cost	$1,907,157	$5,122
Gross discharge income	$4,618,323	$13,645
Total net revenue	$2,711,167	$8,523
Arbitrage revenue per MW nameplate	$9,513	$7,679
Avg. charge price ($/MWh)	$37.98	$43.20
Avg. discharge price ($/MWh)	$138.14	$149.98
Avg. spread ($/MWh)	$100.16	$106.77
Time weighted avg. QLD spot price	$53.96	
Charge price % below QLD spot price	-29.6%	-19.9%
Discharge price % above QLD spot price	156.0%	177.9%
Capacity factor (discharge only)	5.4%	3.8%
Capacity factor (charge + discharge)	14.3%	8.3%
Cap contract % volume cover	74.9%	57.9%
Cap contract % financial cover	73.8%	52.9%

As seen in Table 3.4, the battery's performance was surprisingly close to Wivenhoe's on a range of key metrics, including average spread achieved, and arbitrage revenue when adjusted to a $ per megawatt nameplate basis (and accounting for Wivenhoe's higher capacity factor driven by larger storage volumes). As expected, the battery significantly outperformed Wivenhoe in terms of round-trip efficiency.

The biggest difference observed was between each asset's effectiveness at providing cap coverage. Wivenhoe achieved volumetric and financial coverage percentages in the mid-seventies range, compared to the battery's mid-fifties. This is likely a reflection of Wivenhoe's vastly different storage duration characteristics, as well as the benefits of manual trading by experts as opposed to the shortcomings of DRE's automatic forecast-based trading that are discussed further in section 4.3.


TESLA
TESLA

4. Challenges & Learnings

4.1 Available Capacity and Uptime

Throughout Q1, the availability and capacity of the battery was impacted by two separate challenges.

Firstly, through the commissioning phase of the project a fault with one of the ten battery packs was discovered. As a result, 215kWh of energy capacity was taken out of service. Due to the Powerpack's modular nature, the remaining 90% of system capacity remained available. This outage persisted throughout the duration of the quarter and required spare parts to be obtained for rectification. With the restoration of the faulty module, the optimisation constraints of DRE will be relaxed, thus broadening the control action space for the battery to hedge against spot price volatility.

The second challenge encountered through the first quarter of operation related to sporadic communications dropouts between the battery control system and site level power monitoring equipment. As with most large-scale behind-the-meter energy generation, the battery is required to comply with export restrictions as stipulated by the network service provider. Modbus TCP communications between the battery controller and site level electrical metering are essential in achieving export constraints; without live feedback, no control of the battery can take place. To handle a loss of communications from site level monitoring, physical switching of the LV circuit breakers to both battery inverters was enforced. The operational implication of this meant that upon any loss of communications (no matter how brief) the battery would be electrically isolated from the campus electrical network until an operator manually re-closed the breakers. Through two consecutive weekends in February (8th to 9th and 15th to 16th), a loss of communications for less than one minute caused the tripping of these circuit breakers, taking the battery out of operation until UQ staff manually reset them on Monday morning. These outages contributed almost entirely towards the 5.7% downtime experienced throughout the quarter. As a result of this, the strategy for managing a communications loss was re-engineered so that instead of physically isolating the battery from the grid, control set points are overridden to zero in the battery controller, causing the battery to remain switched on but in standby until healthy communications is re-established (usually only a minute or so later).

4.2 Communications Problems

At 12:09pm on the 23 January, internet connectivity issues within Amazon Web Services (AWS) Sydney (ap-southeast-2) resulted in a loss of connectivity between NEM data sources and DRE. As a result, no new market data (including dispatch pricing and pre-dispatch forecasts) was ingested into the system. Furthermore, connectivity issues between AWS Lambda and EC2 hosted databases resulted in live data being unobservable through the primary data APIs. As a result of this event, no new scheduling optimisations were executed, causing the receding horizon control methodology to break. In this scenario, DRE has been designed to utilise and execute the last generated schedule as the control policy for the battery. While this approach ensures that operation of the battery is not immediately halted, the success of the control is dependent on the quality of the last available forecast. As such, it became the responsibility of UQ staff alerted to the outage to review the operational profile suggested and intervene if required.

By coincidence, this AWS outage happened to also occur on a day of volatile spot pricing in the NEM, with significant ramifications for the battery's Q1 performance. On this day, the last policy generated by DRE before the AWS outage was characterised by a single continuous discharge from 3:30pm to 5:10pm, aligning with when the highest priced trading intervals had been expected to occur. By 3:00pm, the price forecast had evolved, with the peak price period now expected to commence later in the day from 4:30pm. As a result, manual intervention took place and the battery was discharged continuously from 4:35pm to 6:20pm. In the end and with the benefit of hindsight, it can now be seen that the volatility in prices over the evening peak occurred mainly over the trading intervals ending 5:30pm and 7:00pm. Therefore while manual intervention allowed the battery to respond appropriately to the first price spike, the battery had run flat prior to the second peak.

In order to retrospectively review the impact of the AWS outage, DRE's simulation tool was utilised. The simulator makes use of historic forecasts and dispatch prices in order to determine how the battery would have behaved across the afternoon had normal operation occurred. As illustrated by Figure 4.1, had communications been functional, DRE would have succeeded in fully discharging across both peak pricing intervals, almost doubling the resulting daily arbitrage revenue from $1,312 to $2,495.

Figure 4.1: Comparison of actual vs. hypothetical battery performance during AWS outage on 23 January

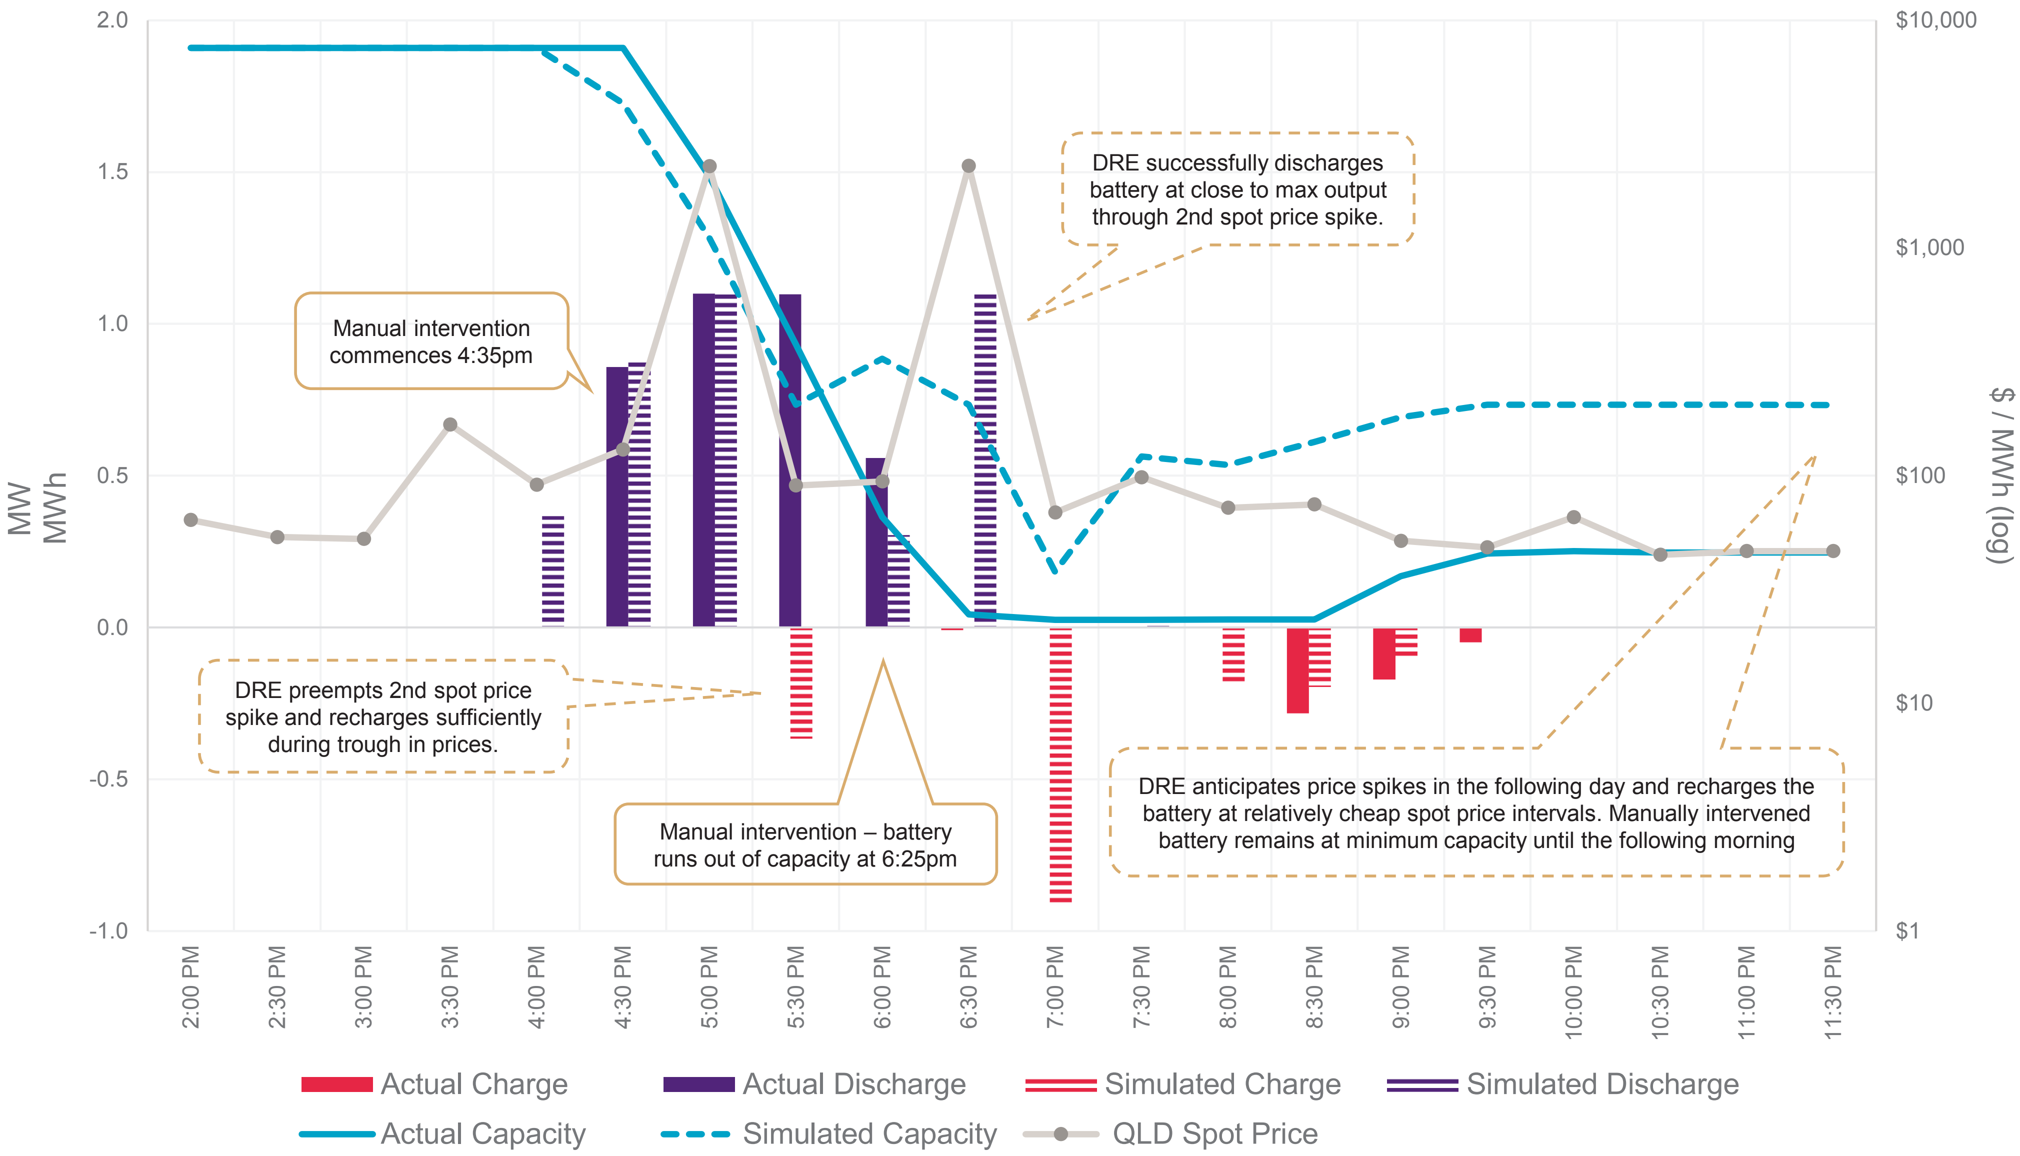

4.3 Imperfect Forecasts and the “One in the Hand, Two in the Bush” Problem

As DRE is a market driven optimisation system it is dependent on the accuracy of price forecasts for that market. As can be expected, any financial market will inherently involve complex dynamics and as such the quality of forecasts will be variable. This is particularly the case in periods of tight supply and demand in the NEM, as well as through unpredictable events such as generator or network outages.

Figure 4.2 illustrates a sliding window snapshot of the policies generated by DRE across the afternoon of January 31. This day was selected as an example of the challenges faced by DRE as the NEM experienced both of the aforementioned drivers of market forecast uncertainty. The supplementary analysis in Table 4.1 provides insight into how this forecast uncertainty impacted the overall control of the battery across the day.

The events of 31 January also allude to a wider challenge of trying to optimise arbitrage in a market with significant uncertainty of price forecasts. It arises from the battery being required to make a choice between charging or discharging during an interval at a relatively certain price or waiting to do so at a later point in time with forecast pricing that may be more lucrative but which has a higher degree of uncertainty associated with it. This has been dubbed by the project team as the “one in the hand versus two in the bush” dilemma. An example would be a thirty minute interval where the battery had a high degree of confidence of being able to achieve a discharge price around $2,500/MWh based on the five-minute dispatch intervals that had already elapsed as well as the P5 dispatch forecast for the remainder of the interval. Objectively, discharging at this price would be a success for the battery. However, in the event that DRE was seeing a price forecast of $14,000/MWh several hours later in the day, it is possible that no energy would be discharged during the $2,500/MWh interval in order to preserve battery capacity. Indeed, as occurred on 31 January it is possible that the battery may even charge at a relatively high price in order to prepare for a yet higher forecast price in the future. As a volatile market with a high degree of price uncertainty, there is every chance that the forecast $14,000/MWh in this example may not eventuate and the battery would have missed the day’s highest price interval. This encapsulates the dilemma of needing to trade off certainty in the moment (“one in the hand”) with the chance of a higher reward in the future that carries with it more uncertainty (“two in the bush”). This dilemma is particularly pronounced with an asset that has relatively limited storage duration such as a battery.

From assessing events such as those on 31 January, it became clear that while MPC is an effective methodology for arbitrage control, in order to mitigate the most severe financial risks associated with inaccurate forecasting, a hybridised control approach should be considered for specific operational conditions. The chosen hybrid control method implemented in the DRE controls thus far revolves around trickle charging the battery during certain rule-based pricing scenarios. How this works in practice is demonstrated with the following examples. This approach is currently implemented for battery charging and is under further consideration for implementation for discharging as well in order to help address the above described “one in the hand, two in the bush” dilemma.

Scenario 1 (business-as-usual)

The battery system is currently on standby and is at its minimum energy capacity of 185kWh (reserved to accommodate FCAS commitments). DRE observes a very mild forecast with prices only ranging between $35/MWh and $70/MWh across the forecast horizon. The MPC controller within DRE sees no arbitrage opportunity due to an insufficient minimum spread and leaves the system in a standby state. Later that day an unexpected network fault in the NEM causes the spot price to suddenly spike to $14,000/MWh. DRE observes this high spot price but the battery is constrained by its low energy state and cannot discharge. UQ wears the full financial impact of the spot price spike.

In this scenario, a pure MPC approach to controlling the battery is misdirected by the mild prices observed in the visible forecast. While this approach avoids charging the battery at energy prices that may not be ‘low’ compared to what could be achieved if an arbitrage spread was being seen in the forecast, the cost of not being able to respond to unexpected spikes in spot price carries a much higher potential risk. This example illustrates the case for needing an approach such as the trickle charge hybrid control method.

Scenario 2 (with trickle charge)

The battery system is currently on standby and is at its minimum energy capacity of 185kWh. DRE observes a very mild price forecast with prices ranging between $35/MWh and $70/MWh across the forecast horizon. Given that the current spot price is $40/MWh and is in the lower range of expected spot price outcomes for the quarter, DRE uses a proportional controller to charge the battery up to 740kWh. This approach uses a trigger price (e.g. $45/MWh) to commence charging at a rate of 300kW, proportionately ramping this all the way up to 1.1 MW if prices hit $0/MWh or below. If at any point an arbitrage opportunity emerges in the forecast, it would supersede this control mode. The following half hour, an unexpected network fault in the NEM causes the spot price to suddenly spike to $14,000/MWh. DRE observes this high spot price and is able to fully discharge the battery over the full 30-minute trading period as the trickle charge hybrid control method ensured that a minimum viable quantity of energy was available to respond to unexpected events.

Continual refinement of this methodology and relevant price thresholds is important to ensure that it does not inadvertently create perverse outcomes. As an example, it was observed during one period in late Q1 that DRE was unable to take advantage of unforeseen negative price intervals as it had fully charged using this trickle charge methodology. Following this incident price thresholds were reviewed (to also reflect a change in underlying NEM price ranges that were being observed), as well as rules put in place to limit the amount of trickle charging to cap out once half an hour of discharge at full capacity was stored (after accounting for the volume already reserved for FCAS).

Figure 4.2: DRE scheduled average MW and pricing forecasts across 31 January

Time	Schedule												Actual		Capacity
	11:00am		12:30pm		2:00pm		3:30pm		5:00pm		5:30pm		Actual		Capacity
	MW	$/MWh	MW	$/MWh	MW	$/MWh	MW	$/MWh	MW	$/MWh	MW	$/MWh	MW	$/MWh	MWh
12:00 PM	■ -	$74.71	■ -										▲ 1.099	$69.18	1.33
12:30 PM	■ -	$77.70	■ -	$77.34									▲ 0.271	$74.62	1.78
1:00 PM	■ -	$78.74	■ -	$90.48									▼ -0.074	$71.67	1.89
1:30 PM	■ -	$82.99	■ -	$92.51									▲ 0.244	$69.93	1.87
2:00 PM	■ -	$82.99	■ -	$94.99	■ -	$79.48							▼ -0.019	-$121.06	1.95
2:30 PM	■ -	$67.01	■ -	$70.99	▼ -0.025	$72.85							▼ -1.047	$117.82	1.94
3:00 PM	■ -	$58.64	■ -	$74.71	■ -	$74.71							▼ -0.740	$119.83	1.42
3:30 PM	■ -	$67.70	■ -	$74.71	■ -	$82.99	▼ -1.005	$99.10					▲ 0.224	$277.89	1.04
4:00 PM	■ -	$70.99	■ -	$74.71	■ -	$74.71	▼ -1.110	$89.85					▼ -0.994	$1,569.42	1.14
4:30 PM	■ -	$82.99	■ -	$91.35	■ -	$94.99	■ -	$591.42					▼ -0.531	$866.89	0.63
5:00 PM	■ -	$91.78	■ -	$91.78	▲ 0.195	$143.32	▲ 0.369	$591.42	▼ -0.930	$94.75			▲ 0.185	$1,609.87	0.36
5:30 PM	■ -	$94.99	■ -	$94.99	▲ 1.110	$160.32	▲ 1.056	$1,494.69	▲ 1.110	$1,580.57	▲ 0.47	$517.99	▼ -0.328	$463.42	0.43
6:00 PM	■ -	$94.99	■ -	$77.65	▲ 1.110	$143.32	▲ 1.110	$1,494.69	■ -	$86.46	■ -	$113.61	▼ -0.111	$898.48	0.27
6:30 PM	■ -	$83.74	■ -	$80.72	▲ 1.110	$201.00	▲ 0.991	$1,494.69	■ -	$82.99	■ -	$592.42	▼ -0.037	$1,162.85	0.21
7:00 PM	■ -	$88.78	■ -	$82.99	■ -	$94.99	■ -	$318.87	▼ -0.181	$70.99	■ -	$90.75	▲ 1.052	$26.50	0.19
7:30 PM	■ -	$78.73	■ -	$78.71	■ -	$82.99	■ -	$82.99	▼ -1.110	$60.89	■ -	$70.99	▲ 0.286	$118.41	0.64
8:00 PM	■ -	$82.99	■ -	$82.99	▲ 0.001	$82.99	■ -	$264.46	▲ 1.110	$262.49	■ -	$266.18	▼ -0.853	$136.09	0.73
8:30 PM	■ -	$78.73	■ -	$78.73	■ -	$70.99	■ -	$136.43	■ -	$97.07	■ -	$123.79	▲ 0.813	$83.06	0.29
9:00 PM	■ -	$74.71	■ -	$74.71	■ -	$68.73	■ -	$94.99	■ -	$94.99	■ -	$108.20	▼ -0.885	$90.96	0.63
9:30 PM	■ -	$70.54	■ -	$70.99	▼ -1.110	$60.89	■ -	$72.99	■ -	$63.33	■ -	$67.01	▲ 0.024	$67.91	0.18
10:00 PM	■ -	$90.89	■ -	$92.72	■ 0.955	$74.71	■ -	$94.99	■ -	$78.73	■ -	$82.98	▲ 0.013	$97.24	0.19
10:30 PM	■ -	$65.09	■ -	$66.08	■ -	$66.96	■ -	$78.73	▲ 0.001	$67.97	■ -	$68.22	▲ 0.009	$66.23	0.19
11:00 PM	■ -	$60.79	■ -	$60.79	▼ -0.001	$62.52	■ -	$73.37	■ -	$66.84	■ -	$66.84	▲ 0.011	$64.63	0.19
11:30 PM	■ -	$61.05	■ -	$61.47	■ -	$63.60	■ -	$68.73	■ -	$66.67	■ -	$66.67	▲ 0.001	$59.61	0.19

Table 4.1: DRE policy analysis across 31 January

Time	Explanation of DRE actions
11:00am 12:30pm	Throughout the morning and early afternoon, DRE observed energy prices in the range of $60/MWh to $100MWh. In response to this mild forecast, the resulting policy of the battery was to remain in standby throughout the day.
2:00pm	By 2pm the spread in forecast energy prices widened, introducing arbitrage opportunities. By utilising the current state of the battery along with a system model for the battery, DRE optimises the operational profile over the observed forecast horizon to minimise energy costs. As the battery was almost fully charged by 2:00pm, DRE schedules the battery to discharge over the peak pricing period from 5:00pm to 7:00pm.
3:30pm	At 3:30pm the forecast for the evening peak was predicted to be volatile and sustained with 3 consecutive trading intervals exceeding $1,400/MWh, as well as 2 trading intervals of $591/MWh. Due to instability in the forecast through the previous 90 minutes, DRE had directed the battery to discharge 900kWh which now needed to be recharged in order to prepare for this predicted spike in prices. What can be observed in the resultant operational profile of the battery is that by the conclusion of the 3:30pm interval ($277.89/MWh), the opportunity to pre-emptively recharge the battery had been missed, meaning that only 100kWh of energy was charged back into the battery. Given the rapidly changing price forecasts throughout the mid-afternoon period, the battery went into the peak pricing period with just 1.14MWh or 60% of its expected full charge energy.
5:00pm	By the early evening the QLD energy market had already experienced a $1,569/MWh trading interval through which the battery appropriately discharged at close to its rated limit. What can be seen from the actual outcomes is that the trading interval beginning at 5:00pm would also settle high at $1,609/MWh. Unfortunately, as seen in the forecast received by DRE at the time, the predicted trading price for the interval was $94/MWh, with the interval beginning at 5:30pm expected to peak at $1,580/MWh. As a result, DRE scheduled the battery to charge sufficiently to ensure that full discharge could be achieved through the expected 5:30pm to 6:00pm peak; a scenario which did not eventuate. It can be seen that despite initially planning to charge at 930kW, the resultant profile of the battery shows an average charge of just 185kW. This is due to the fact that as the 30 minute trading interval progressed and the price spiked at 5:20pm, DRE was able to counter the charging that happened at the start of the trading period by discharging through the end of the period. The net effect over the interval however was still one of charging.
5:30pm	With high energy prices being observed in the current trading period ($517.99/MWh), DRE transitions into a more reactive scheduling strategy, allowing it respond to newly dispatched prices as fast as possible.

5. Next Steps

5.1 Demand Forecasting

In order to make intelligent decisions about the optimal usage of the battery while taking into account monthly demand charges, forecast models of the net load for the St Lucia campus need to be developed. The optimal scenario would be to accurately forecast the 30-minute peak demand intervals for each month at the start of the month. This would ensure that the battery could be optimally scheduled to co-optimise revenue from both arbitrage and network charges. In reality though, a 30+ day forecast would carry a significant degree of uncertainty, effectively to the point of being unusable. This is directly attributed to the tight correlation between air conditioning system load (the largest driver of UQ's energy consumption) and ambient weather conditions. As the ability to predict weather on a medium or long-range scale remains a significant challenge, it becomes clear that uncertainty will arise in any demand forecast underpinned by HVAC systems.

Notwithstanding this, by utilising gridded weather forecasting models from the Bureau of Meteorology as well as a range of historic date, time, and campus related load data, machine learning methods have yielded promising results in developing 30-minute resolution, one-week long forecasts for net campus load. While not an ideal scenario (as co-optimisation of peak demand lopping and arbitrage can only occur over a 7-10 day time horizon at most), it is still expected that this will enable UQ to begin realising revenue from peak demand lopping on top of the other services discussed in this report. It is planned for these forecasting models to be developed in a serverless cloud environment and integrated into DRE's control algorithms by Q3 2020.

5.2 Price Uncertainty

One of the key forecasting challenges currently faced by DRE in trying to maximise arbitrage revenue is the inherent unreliability of AEMO's pre-dispatch price forecasts combined with the need to optimise the battery's net position over a half hour settlement period made up of the average of six individual five-minute dispatch intervals. The introduction of five-minute settlement will help to at least partially address this problem, as the battery will reliably know the final price it is able to obtain for charging or discharging as soon as each interval's dispatch price is released by AEMO. This however still leaves the challenge of trying to produce and optimise a schedule for charging and discharging to ensure that the lowest prices of the day are selected to charge, and that enough energy is held in storage for trying to pick the highest prices of the day to discharge.

Exploring opportunities to rely less on the default pre-dispatch price forecasts provided by AEMO is a key area of focus for the project team moving forward. Options currently under consideration include the use of machine learning and other techniques to produce independent price forecasts based on a wide array of market data inputs, or the use of AEMO's existing price forecast sensitivities to look at the range of potential pricing outcomes in each interval as opposed to a single value. In reality, it is likely that the optimal solution will involve a combination of different approaches as well as further refinement of methods like the previously discussed trickle charge hybrid control to ensure a degree of 'rules based' protections against random market outcomes.

It is important to note that there is a delicate balance required when handling forecast uncertainties in predictive control approaches. Where uncertainty models can be utilised to better evaluate expected values in forecast variables, the inclusion of these approaches needs to be carefully balanced against any potential cost to response time. For example, in a wholesale market operating on 5-minute trading intervals, the difference between a 15 second re-evaluation time and a 1-minute re-evaluation time is a significant portion of the trading window. Any delay in the battery's response decision will have a financial impact in the form of a reduced volume of energy being charged or discharged, and minimising response time as far as possible should always be a key goal. These factors will need to be carefully considered as new approaches to managing price uncertainty are trialled with DRE.

5.3 Demand Response Portfolio Expansion

The power of UQ's decision to build DRE as a custom in-house system lies in its ability to scale up with a portfolio of controllable energy assets beyond just traditional batteries. As an example, work is currently underway to integrate UQ's recently commissioned 3.5 megalitre thermal energy storage tank (Figure 5.1) into DRE. This tank acts as a form of water battery and is able to 'charge' by storing cooling energy that can then be 'discharged' later to offset energy usage associated with HVAC systems. This system can be modelled much like a conventional battery, although additional constraints and modifications to account for thermal efficiency effects and physical limitations like machinery response times will be required.

DRE is also already being expanded to help schedule the operation of electric vehicle charging stations, ensuring that UQ's fleet vehicles are charged to a sufficient level before the start of the workday while also ensuring that charging does not occur during peak energy price trading periods. This expands the role DRE is already playing in being able to temporarily curtail some HVAC loads in response to extreme spot market price outcomes.

The learnings from DRE's implementation on the battery project are critical to continuing the rollout of UQ's overall 'Gensumer' strategy, with there likely to be several megawatts of load and generation being autonomously orchestrated and controlled by the system by the end of 2020. This will also feature the new challenge of optimising UQ's load based on the net exposure to spot prices when taking into account generation from the Warwick Solar Farm, which will be fully operational in Q3 2020.

Figure 5.1: UQ Gatton central energy plant with thermal energy storage

Acronyms

AEMO	Australian Energy Market Operator
AWS	Amazon Web Services
DR	Demand Response
DRE	Demand Response Engine
DUoS	Distribution Use of System
FCAS	Frequency Control Ancillary Services
LGC	Large-scale Generation Certificate
MPC	Model Predictive Control
NEM	National Electricity Market
SCADA	Supervisory Control and Data Acquisition
STC	Small-scale Technology Shortfall
TUoS	Transmission Use of System
UQ	University of Queensland

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

CREATE CHANGE