| DOCKETED | |
|---|---|
| **Docket Number:** | 21-DR-01 |
| **Project Title:** | Supply Side Demand Response |
| **TN #:** | 246244 |
| **Document Title:** | CEC Proposal Hourly Regression Capacity Counting Methodology for Supply-Side DR |
| **Description:** | N/A |
| **Filer:** | Courtney Wagner |
| **Organization:** | California Energy Commission |
| **Submitter Role:** | Commission Staff |
| **Submission Date:** | 9/28/2022 2:30:20 PM |
| **Docketed Date:** | 9/28/2022 |

California Energy Commission

**DELIBERATIVE REPORT**

# Hourly Regression Capacity Counting Methodology for Supply-Side Demand Response

## CEC Working Group Proposal

**September 2022 | CEC-XXX-XXXX-XXX**

# California Energy Commission

Erik Lyon

**Primary Author**

Tom Flynn

**Project Manager**

# TABLE OF CONTENTS

Page



# LIST OF FIGURES

Page

# LIST OF TABLES

# EXECUTIVE SUMMARY

The hourly regression capacity counting methodology proposal is centered on a recurring cycle of ex-post capacity measurement and ex-ante capacity projection. The ex-post measurement is highly standardized such that all parties can agree in advance to the measurement procedure and outcome; the ex-ante projection is much more flexible to give DR providers the ability to account for expected changes in their resources such as enrollment and customer composition. Figure 1 summarizes this cycle.

**Figure 1. Ex-post and ex-ante capacity cycle**

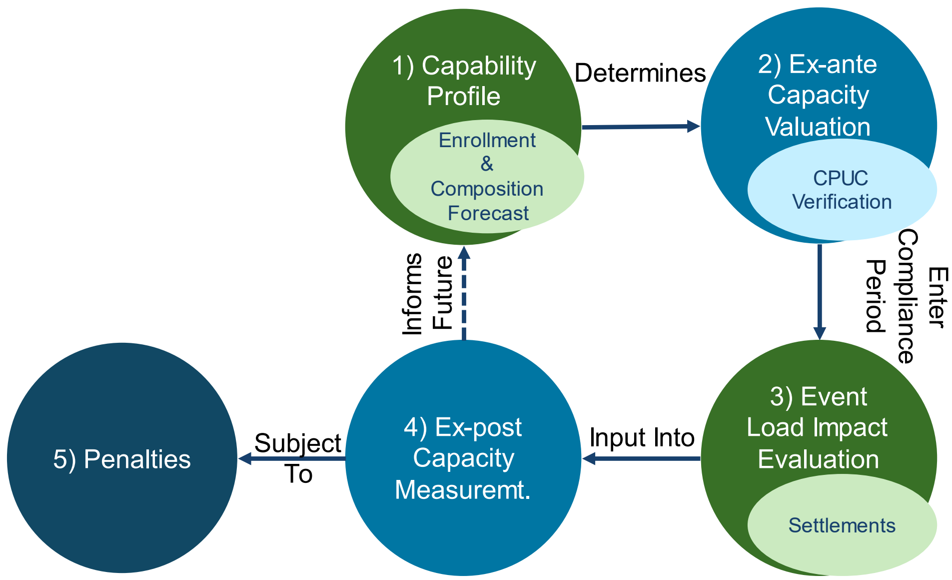


Each step in the process in additional detail below.

1. **The DR provider creates a capability profile for its resource or aggregation of resources for each hour.** The capability profile is a projection of how the resource can be expected to perform under varying temperature conditions for every hour. A profile is required for each combination of month and hour slice for which a capacity value is sought. Figure 2 shows a stylized version of a capability profile for a single hour, which can be applied to one or more months. DR providers may submit any of the points A–D to define the temperature sensitivity of a resource, but all are optional.

**Figure 2. Diagram of Ex-ante Hourly Inputs**

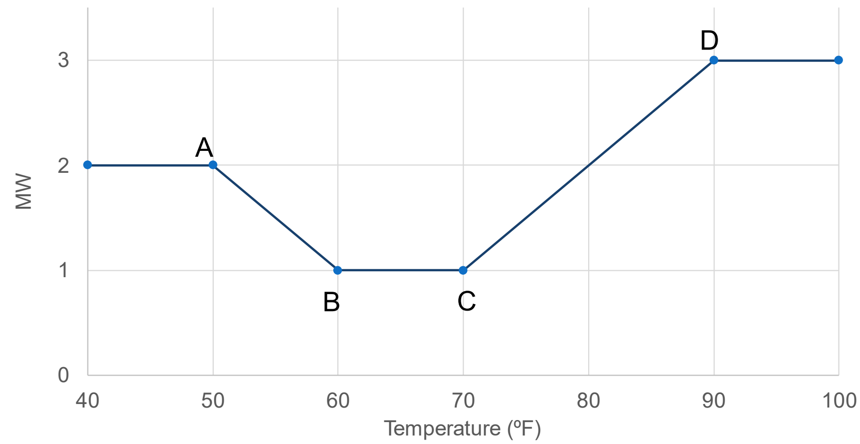


In general, this capability profile will be informed by the previous years' ex-post model and capacity measurement (step 4) but can be changed by the DR provider to take any changes in the resource such as enrollment or customer composition into account, or control for issues resulting in underperformance that have been resolved. New resources will be required to submit a capability profile as well.

2. **The capability profile directly determines the ex-ante capacity value of the resource, which is subject to a finding of reasonableness by CPUC Energy Division staff.** The ex-ante capacity is defined as the value of the capability profile at the planning temperature. Figure 3 shows a graphical representation of ex-ante capacity determination for a resource that shows sensitivity only to high temperatures (for example, a resource targeting air conditioning) by submitting only points C and D in the capability profile (step 1).

**Figure 3. Graphical Illustration of Ex-Ante Capacity Determination**

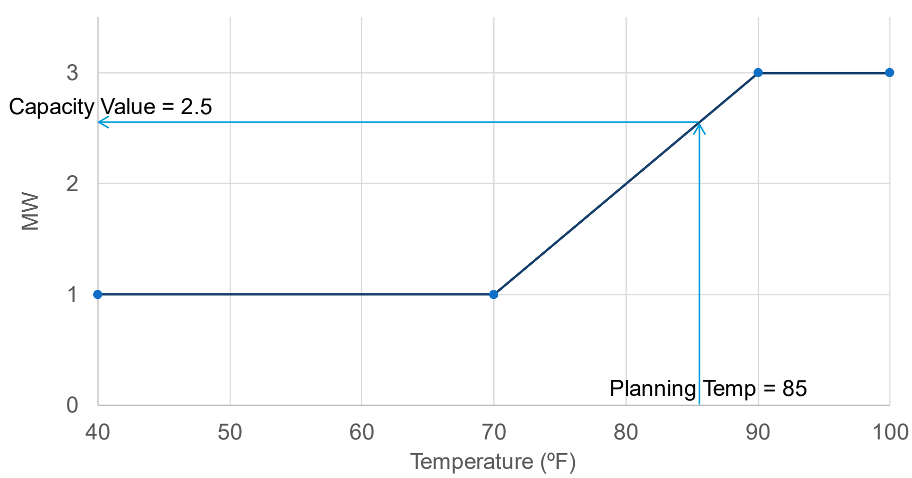


All requested capacity values no greater than 25 percent above the previous year's ex-post capacity measurement (step 4) shall be granted, so long as the resource met at least 90 percent of its committed capacity in the previous year. Resources requesting an increase of over 25 percent or that delivered less than 90 percent of committed capacity in the previous year are subject to more detailed staff review.

3. **Individual ex-post load impacts are calculated.** When possible, it is preferable to use the same baseline methodology used to calculate settlements in the California ISO energy market for consistency between the operational and planning space. However, an alternative method may be used when two conditions are met: 1) the alternative method can be shown to be more accurate than the settlement method, and 2) the alternative method is infeasible to implement for settlement. These load impacts will be adjusted relative to the amount bid according to the following equation:

$$\text{BNLI} = \max\left(\text{Bid}\left(\frac{\min(\text{Delivered, Dispatch})}{\text{Dispatch}}\right), \text{Delivered}\right)$$

4. **Ex-post capacity is determined.** Using individual load impacts (step 3) and any changepoints submitted in the ex-ante capability profile (step 1), a linear regression model of ex-post demonstrated capability is developed, which in turn determines the capacity value.

5. **Penalties are applied to any shortfall in delivered capacity.** Penalties are assessed on the portion of capacity to which the resource was committed but failed to demonstrate. Relative to the *committed* capacity (defined as the lesser of contracted capacity and QC), the penalty is equal to two times the shortfall below a minimum performance threshold of 94.5 percent. That is, the resource shall be compensated for its *delivered* capacity minus again the shortfall. Figure 4 shows the revenue a DR provider will receive as a proportion of its contract value as a function of the demonstrated capacity it delivers ex-post. DR providers may include contract provisions for compensation above the minimum contracted capacity, as shown up to 110 percent below.

**Figure 4. Effective Capacity Revenue under Capacity Shortfall Penalty**

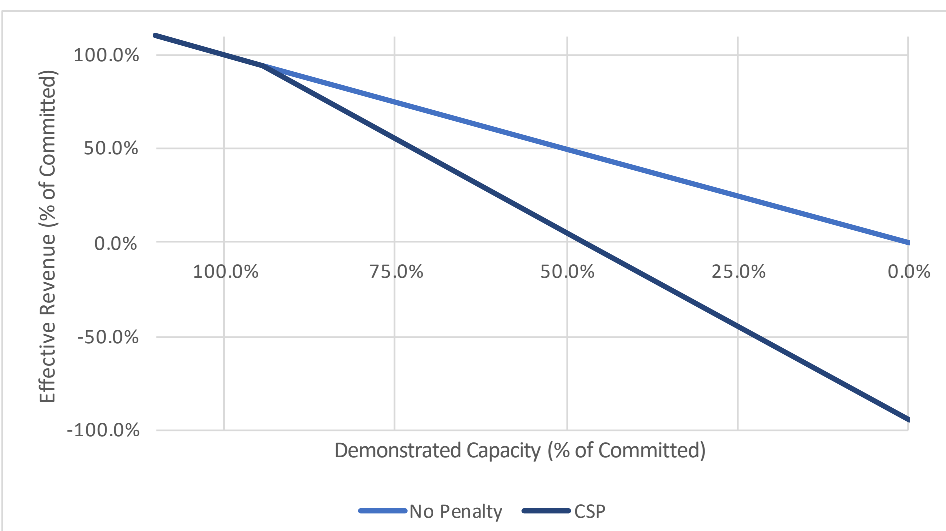

This report is prepared for consideration by members of the CEC working group on the qualifying capacity (QC) of demand response (DR) resources, which includes representatives from DR and storage providers, evaluation consultants, utilities, as well from CEC, CPUC, and the California ISO. The report has not been formally published by the CEC and the views expressed are solely those of its author.

# CHAPTER 1:
# Inputs for Capacity Calculation

Building a capacity calculation model requires defining the inputs. This chapter introduces constraints on customer energy baseline for calculating individual load impacts, and how individual load impacts with be adjusted for partial dispatches. These inputs are used in step 3 of the overall process. They are described first in the report because they are foundational to the overall cycle.

## Customer Energy Baselines

To measure delivered load impacts from individual dispatches relative to a counterfactual baseline, we propose a simple rule. The baseline method used to measure load impacts for the purpose of calculated ex-post capacity shall be the same baseline method used to calculate settlements in the California ISO energy market unless two conditions are met:

1) The alternative method can be shown to be more accurate than the settlement method.
2) The alternative method is infeasible to implement for settlement.

This above rule aims for consistency between the operational and planning space when possible, but allows flexibility when required.

### Baselines for Weather-Sensitive Resources

One such baseline approach that would likely meet the above conditions is the use of comparison groups. DR providers have long noted that the available methods for determining DR participants' counterfactual settlement baselines are inaccurate for weather-sensitive DR resources such as air-conditioning cycling and similar programs.[1] (Sufficient baselines already exist and are in common usage for non-weather-sensitive resources.) CEC staff finds that the absence of an accurate baseline that employs transparent methods that are fixed and agreed to *before* measurement undermines policymakers' confidence in DR and is a significant barrier to allowing the DR market to reach its full potential.

Such a method has recently been tested, validated, and affirmed as tariff-compliant for the California ISO.[2] This method, a type of comparison group, satisfies the CEC's conditions above for a weather-sensitive settlement baseline.

---

1 Duesterberg, Matt. *Deep dive into OhmConnect's community response during Summer 2020.* OhmConnect. January 13, 2021. https://www.ohmconnect.com/thought-leadership/deep-dive-into-ohmconnects-community-response-during-summer-2020.

2 Glass, Joe, Stephen Suffian, Adam Scheer, and Carmen Best. Prepared by Recurve for the California ISO. November 2021. http://www.caiso.com/Documents/DemandResponseAdvancedMeasurementMethodology.pdf.

However, barriers remain to successfully implementing these baseline methods because they rely on access to nonparticipant hourly electric meter data, to which DR providers (particularly third-party providers) often do not have access. The CEC currently collects this data from the IOUs and large publicly owned utilities and intends to be the energy data hub for California. As such, CEC staff propose the CEC investigate how to take on the role of developing tariff-approved comparison group baselines for providers of weather-sensitive DR resources.

Currently, CEC plans to receive data from utilities approximately quarterly. This frequency is insufficient to clear energy market settlements on a near-daily basis. However, it may be sufficient for the purposes of calculating ex-post delivered load impacts and capacity for performance verification and RA compliance purposes.

## Normalizing Load Impacts for Availability

We propose a measure of bid-normalized load impacts that a hybrid of bid, dispatch (or test), and load impact data. Bid-Normalized Load Impact (BNLI) is calculated according to the following formula for any period in which a DR resource receives a dispatch, including a partial dispatch:

$$\text{BNLI} = \max\left(\text{Bid}\left(\frac{\min(\text{Delivered, Dispatch})}{\text{Dispatch}}\right), \text{Delivered}\right)$$

Intervals in which a DR resource has RA obligations but does not bid will be assigned a BNLI of zero.

Table 1 illustrates the proposed definition of bid-normalized load impact over different scenarios. Under a full dispatch (example 1), the BNLI is equal to the delivered load impacts. Under a partial dispatch, the bid amount is adjusted by the ratio of delivered load impacts to the bid amount (example 2 and 3), but this ratio is always capped at 1 by the minimum function, limiting BNLI to the bid amount (example 4). The only time BNLI can exceed the bid is when load impacts exceed the bid, regardless of the dispatch amount (examples 5 and 6).

**Table 1. Bid-normalized Load Impact Examples**

| Example # | Bid | Dispatch | Delivered | BNLI |
|---|---|---|---|---|
| 1 | 100 | 100 | 90 | 90 |
| 2 | 100 | 60 | 30 | 50 |
| 3 | 100 | 60 | 60 | 100 |
| 4 | 100 | 60 | 80 | 100 |
| 5 | 100 | 100 | 120 | 120 |

| Example # | Bid | Dispatch | Delivered | BNLI |
|---|---|---|---|---|
| **6** | 100 | 80 | 120 | **120** |
| **7** | [Test] 100 | [Test] 100 | 120 | **120** |

DR tests can be used in the absence of ISO dispatches if necessary. In these cases, the amount bid and dispatched should be assumed to equal the entire resource. That is, the concept of a partial dispatch is not applicable to a test event. Mathematically, this implies that the amount that the DR provider believes it can provide (the "bid") and the amount it is attempting to provide (the "dispatch") are the same and these two quantities cancel out in the formula, and the result is simply the delivered load impact (example 7). In the case of a test event, the DR provider does not need to include bid or dispatch values, but the result is conceptually compatible with actual dispatch data.

The hourly ex-post capacity valuation model takes these BNLI values along with the corresponding temperatures as inputs. The recommended granularity is by sub-LAP, but the proposal could be modified to the level of granularity needed, such as by IOU service territory

# CHAPTER 2: Ex-post Hourly Capacity Counting Methodology

This chapter details the proposed ex-post hourly regression capacity counting methodology, step 4 of the overall process. While the ex-ante determination precedes the ex-post calculation chronologically, this proposal is anchored on a consistent, transparent ex-post measurement methodology and so is presented first.

For each hour in each month, the methodology includes the following steps, which are described in greater detail in the following sections:

1. **Run a regression of availability as a function of temperature:** Create a linear regression model of adjusted load impacts on temperature over each month (or grouping thereof) by hour of day. The measure of temperature may include predetermined change points if necessary to account for diminishing resource capabilities under extreme conditions. The regression line generated by this model is the resource's ex-post hourly capacity profile.
2. **Determine the hourly capacity value by the intersection of the availability profile with the monthly planning temperature.** Apply the planning temperature to the capacity profile function to generate an estimate of the resource's capacity value under planning conditions.

The methodology is designed to accommodate weather-sensitive resources but can be simplified further for non-weather-sensitive resources: simply take the average of adjusted load impacts by hour within each month. However, there a standardized definition of weather sensitive or non-weather sensitive is not required; any resource may use either pathway.

## Availability Regression

The regression of DR adjusted load impacts on temperature serves to account for the capacity value of the resource under planning temperatures. The regression may include change points as necessary to account for DR resources with capabilities that change under warmer and cooler conditions, as well as diminishing capabilities under more extreme conditions if necessary. The changepoints must be selected by the DR provider ahead of the RA month in question. The availability regression (and associated changepoints) is not required for non-weather sensitive resources.

Consider a four-hour weather sensitive economic DR resource (PDR) with takeback in the two hours before and after event. On the "worst day," the grid need of the LSE is from 5:00–9:00 p.m. Figure 5 shows the per-meter availability of the resource from the four dispatch hours as well as the four hours with takeback. The light blue dots show the BNLI on each day, and the dark blue lines show the regression results. The DR provider has provided limiting changepoints in September for all dispatch hours, but in the 6:00 and 8:00 p.m. hours of the month the temperature never hit the changepoints during a dispatch.

**Figure 5. Hourly September Availability, 3:00–11:00 p.m.**

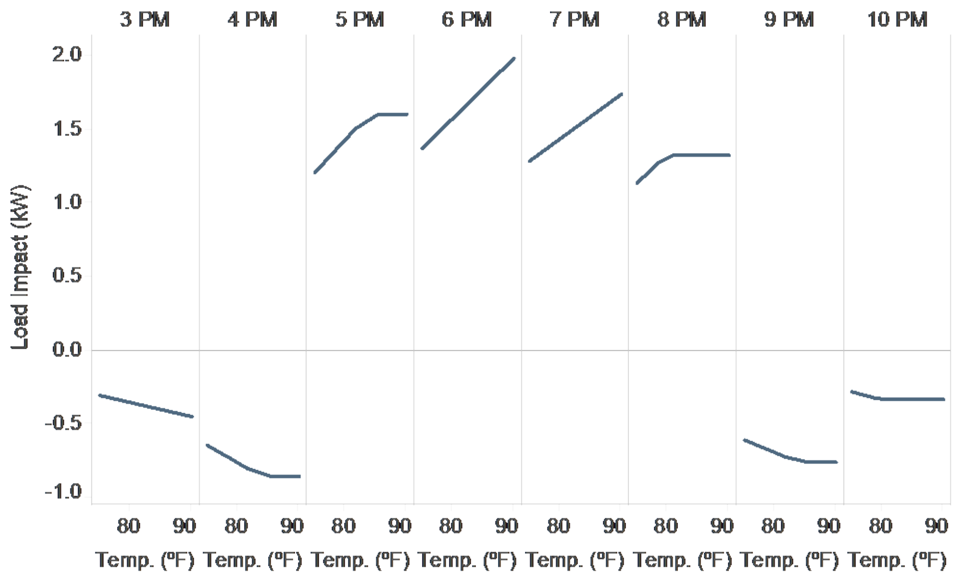


The slope of the regression lines indicates this resource is highly weather sensitive with greater positive and negative values on hotter days, both in the load impacts and in takeback.

## Capacity Value by Hour and Month

The second step is to apply the monthly planning temperatures to the regression line to determine the capacity value in each hour, including takeback hours. Figure 6 shows the same resource as in Figure 5 with planning temperature (vertical dashed lines at 81.3ºF). Where the availability profiles intersect the planning temperature are the hourly September capacity values (horizontal dashed lines).

**Figure 6. Hourly September Availability with Planning Temperature and Capacity Value Superimposed**

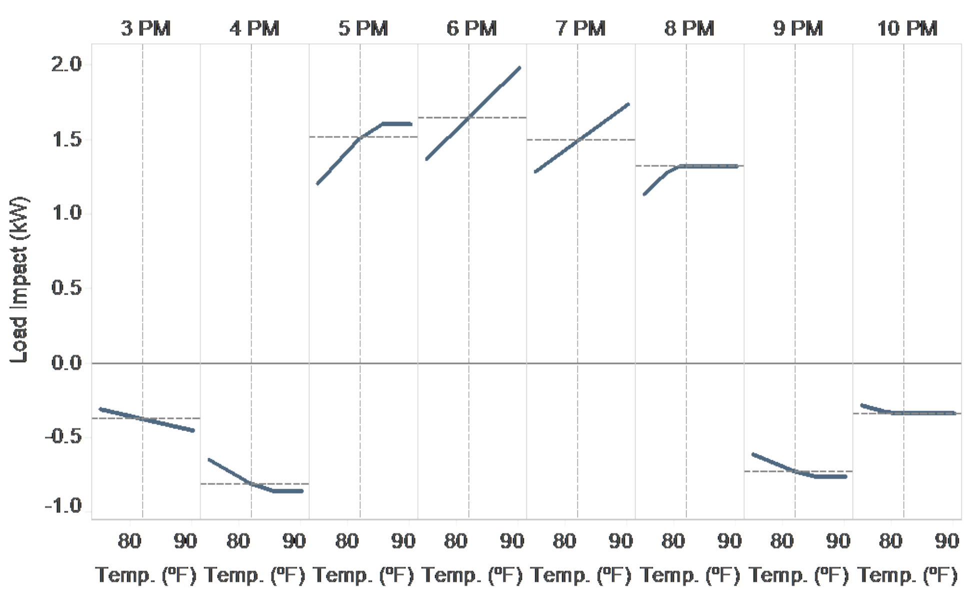

The resulting September capacity values are summarized by hour in Figure 7.

**Figure 7. Hourly August Capacity Value Summary**

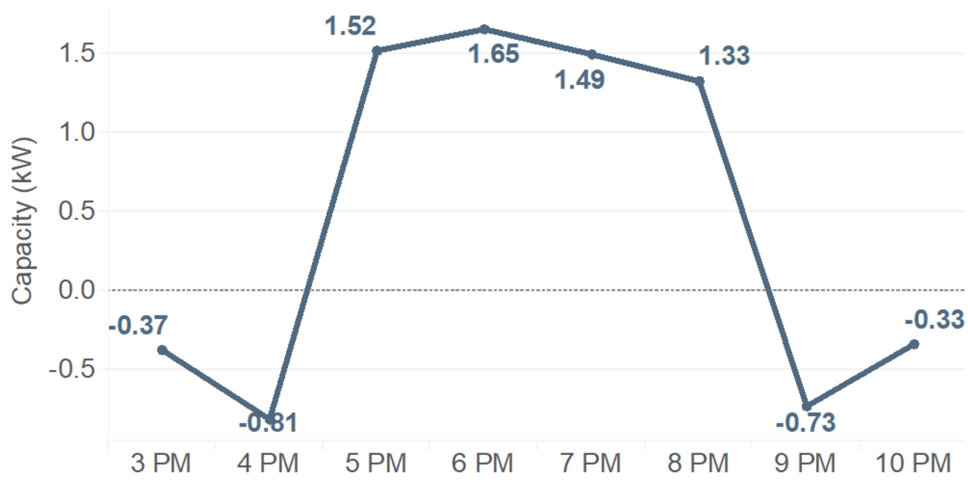


## Additional Considerations

The example above is a simple representative illustration of this methodology. However, real-world circumstances may be more complicated. This section addresses some of these possibilities, as well as minimum dispatch requirements to satisfy the model.

### Winter Months

During winter months (defined here as December through March), peak net load and wholesale prices tend to increase with lower temperatures. Accordingly, the 1-in-2 minimum temperature is used during these months rather than the maximum under this proposal. Figure 8 shows the same process for developing December capacity values (labeled).

**Figure 8. December DR Availability, Availability Profile, Planning Temperature, and Capacity Value**

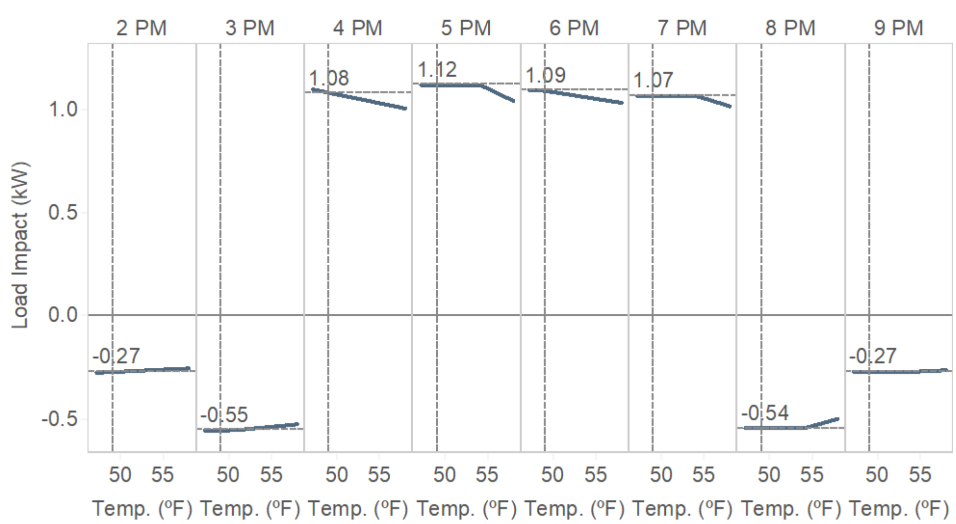

## Economic Bidding

A possible challenge to a 24-slice RA paradigm is the divergence between grid needs under planning conditions and in operations. For example, a DR resource might commit to 5:00–9:00 p.m. availability based on the grid needs of the "worst day," but wholesale prices may be higher from 4:00–8:00 p.m. or 6:00–10:00 p.m. on a given day. For an economically efficient dispatch, the DR resource should be able to shift its bid window without jeopardizing its capacity valuation. Consider the case of an equivalent DR resource to that shown previously except that it bids into the four highest price hours per day rather than a fixed four-hour window. Figure 9 shows the hourly availability of this resource for August.

**Figure 9. Hourly August Availability with Economic Bidding Behavior**

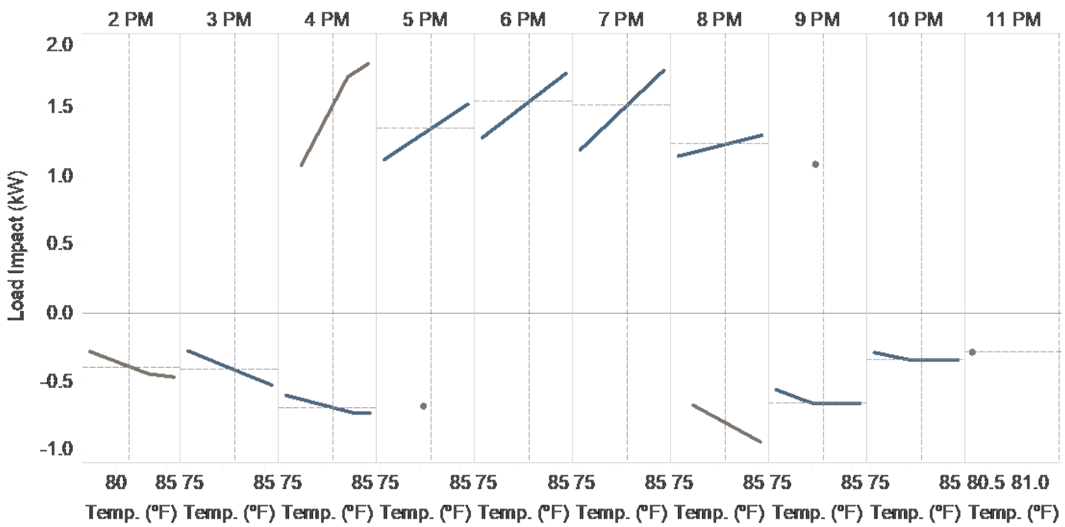


To illustrate, note the 4:00 and 8:00 p.m. hours, both of which include dispatches and takeback. In those hours, the regression handles negative (takeback) and positive (load impacts/reductions) separately, so that only dispatches are considered in hours with RA commitments and only takeback is considered in hours without. As a result, the process of measuring the value of capacity regression in the committed hours produces negative load impacts in the 4:00 hour and positive values in the 8:00 hour, consistent with how the resource is expected to perform on the hottest days when reliability concerns are greatest.

Because the capacity value is measured on the "worst day" that coincides with the highest expected temperature, the methodology only picks up dispatches between 5:00–9:00 p.m. — the same hours as in the fixed window – and takeback in the adjacent hours. However, the hourly capacity values are changed slightly. Figure 10 shows the hourly capacity values for the economic bidding behavior (dark blue) relative to the fixed window (light blue).

**Figure 10. September Hourly Capacity Values for Economic Bidding Relative to Fixed Window**

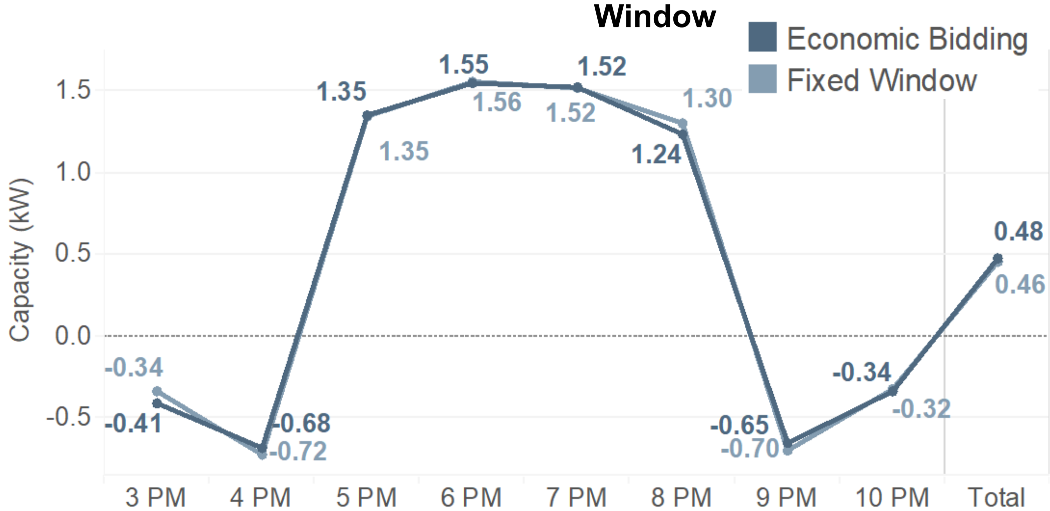


While the positive load impacts are somewhat smaller in the 8:00 hour when the resource bids into the highest-price days, but the takeback in the surrounding hours is also less. These differences generally cancel out one another, and the average capacity value over these hours is slightly *larger* when able to adjust the bidding window. The difference between the average capacity value in the presumed bidding window on the "worst day" for which DR is awarded capacity value and the actual average capacity value may serve as a basis for evaluating performance and assessing penalties. Note that hours with takeback outside the shown hours are discarded and do not influence the final capacity valuation.

## Minimum Dispatch Requirements

In order to successfully generate ex-post capacity measurements using this regression approach, multiple data points are required. This section discusses how the methodology will handle few or zero data points and how this treatment provides an incentive for DR providers to be dispatched in the market. However, there is no specific minimum dispatch requirement.

The ex-post regressions will be run based on the months or "seasons" (defined as any grouping of months defined by the DR provider) in the submitted ex-ante capability profiles described in the following chapter.

In the absence of any dispatch or test results in a season and hour, DR resources will be awarded an ex-post capacity of value of zero. In the case of a single dispatch or test, that single value will be used for capacity across all temperatures.

With a small number of data points, a regression line will still be fit. While such a regression may produce volatile results, that volatility provides an incentive for DR providers to dispatch frequently enough to generate sufficient data to develop a robust ex-post model and support their QC claims. However, there is no specific minimum requirement for what constitutes sufficient data. We also note the ability to combine months into "seasons" can allow DR providers to develop that data set over more months, if the resource behaves consistently relative to temperature across multiple months. This approach will help resources that dispatch less frequently to nonetheless develop evidence of capacity value.

As an extreme example, a non-weather sensitive RDRR resource that dispatches very infrequently may use the entire year as a season. Even if the resource is only called for its two annual test events and never dispatched under emergency conditions, those two test events may form the basis for capacity value for every month of the year.

# CHAPTER 3:
# Ex-Ante Capacity Determination

The ex-ante determination is made based on the DR provider's assessment of its capability to meet a capacity value determined by the ex-post capacity model described in the previous chapter. While the ex-ante process is presented as steps 1 and 2 in the overall cycle, the ex-ante portion relies on knowledge of the ex-post methodology. The process consists of two main steps, analogous to the ex-post process:

1. **Determine the ex-ante capability profile of a resource.** The DR provider submits a stepwise function estimating ex-ante capabilities under varying temperature conditions. Separate capability profiles are required by hour of the day for which the resource has RA commitments. Capability profiles may apply to one or more months.
2. **Determine the ex-ante hourly capacity value by the intersection of the availability profile with the monthly planning temperature.** Apply the planning temperature to the capability profile function to generate an estimate of the resource's capacity value under planning conditions.

## Determining the Ex-Ante Capability Profile

The DR provider will be required to submit parameters defining the underlying capabilities of a resource. At its most basic, the capability profile is what the DR provider forecasts the ex-post model will be. However, the provider may make any adjustments from previous ex-post models to account for growth in enrollment, change in customer characteristics, errors and misfires from previous years that have been resolved, and any other factors deemed necessary by the provider. Critically, the ex-ante capability profile need not be a *predictive* model of the future capabilities of a resource, but a minimum threshold of capabilities that the DR provider can commit to with reasonable confidence. This reframing allows DR providers to adjust capability profiles to include factors like the probability of reaching use limitations such as maximum hours, customer fatigue over multi-day dispatches, and others.

DR providers submit the capability profile by defining parameters for one or more change points that determine a resource's capabilities at different temperatures. Figure 11 shows a schematic of the possible inputs, marked as points A–D. Points B and C represent the points below and above which the resource shows temperature sensitivity. Points A and D represent saturation change points below and above which the resource no longer shows temperature sensitivity.

**Figure 11. Diagram of Ex-ante Hourly Inputs**

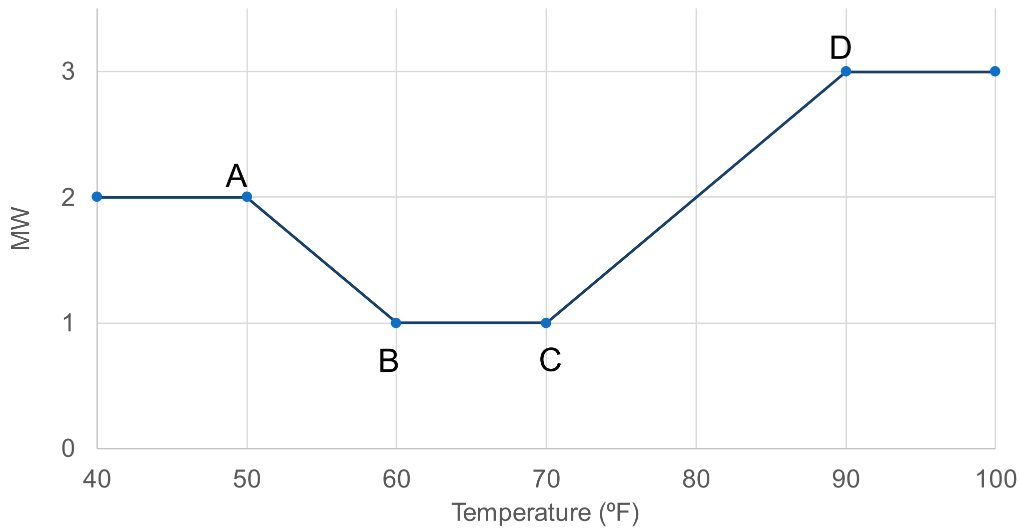


None of these parameters are required, providing DR providers significant flexibility in defining the contours of their resource. For example, a resource targeting heat pumps used for both space heating and cooling might include all four points; a resource targeting cooling-only air conditioners might only require points C and D because it has no cool weather sensitivity.

However, there are a few constraints that must be imposed on these points. These effectively require the profile to behave similarly to the schematic shown above. In other words, points A–D must be in ascending order and points A and D must be higher than points B and C. For hours with takeback, however, the above load impact constraints are inverted to allow for greater takeback under more extreme temperatures (e.g., precooling). Constraining the *absolute value* of load impacts allows the rules to be applied to takeback hours as well. These constraints are summarized in Table 2.

**Table 2. Ex-Ante Profile Change Point Constraints**

| Point | Required? | Temperature Constraint | Load Impact Constraint |
|---|---|---|---|
| A | No | A<B | \|A\|>\|B\| |
| B | Only if A submitted | A<B≤C | B=C |
| C | Only if D submitted | B≤C<D | C=B |
| D | No | D>C | \|D\|>\|C\| |

Capability profiles must be submitted for every hour of every month for which a resource is seeking an hourly QC value, plus any hours in which takeback is expected on the worst day. However, a single capability profile may be used for multiple months and/or hours as appropriate. For example, one profile for a given hour might be used for all months in the summer rate period, and a second would be used for non-summer months. Takeback profiles are optional. However, profiles of zero takeback will be imputed in the two hours before and after QC awards if not provided to ensure DR providers are not simply ignoring takeback.

Critically, the temperature values of the changepoints must be committed to ex-ante and applied when evaluating ex-post. The load impact values will be determined by running a regression of performance relative to temperature, subject to these predetermined changepoints.

## Determining the Ex-Ante Capacity Value

The ex-ante capacity value can be determined unambiguously from the ex-ante capability profile. For each capability profile, the capacity value is the load impact (MW) value that corresponds with the planning temperature for that month. Figure 12 illustrates this graphically: the planning temperature (85ºF) can be traced from the x-axis to the capability profile, and then followed to the intersection with the y-axis to reveal a capacity value of 2.5 MW.

**Figure 12. Graphical Illustration of Ex-Ante Capacity Determination**

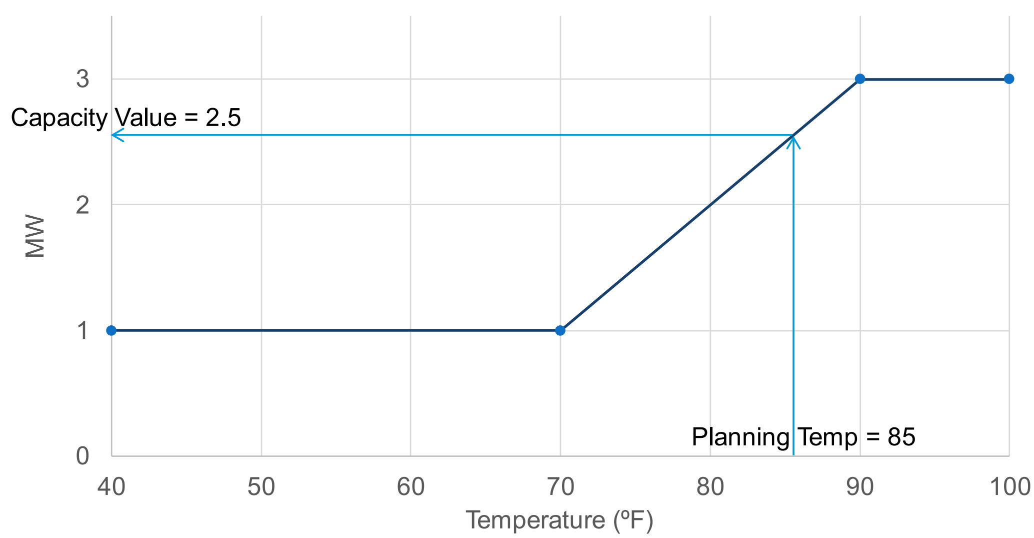


## Review and Approval

CPUC staff retain the role of approving final DR QC values. Other DR QC working group stakeholders have submitted proposals including requirements for reporting requirements, including data and evidence for the capability to meet future capacity obligations.[3] This proposal does not weigh in on the specific reporting requirements, but notes the existing

3 See proposals from OhmConnect and California Energy Efficiency and Demand Management Council for possible reporting requirements.

process includes many reporting requirements not currently required for determining QC in the following year, and we support streamlining these requirements as appropriate.

This proposal includes constraints on how CPUC staff can adjust QC values, however. To adjust the final QC, the underlying capability profile must be changed. This can be done by adjusting the MW values of the change points that result in the desired final QC value while preserving the relationship between the capability profile and the final capacity value.

### Streamlined Approval

We also propose a streamlined approval process for DR providers and resources that have a proven track record and are growing at a reasonable pace year-over-year. Specifically, we propose for CPUC staff automatically approve requested QC of any resource aggregation for any hour and month that meets the following two criteria:

1. Ex-post capacity value is at least 90 percent of the committed capacity.
2. Requested ex-ante capacity is no more than 25 percent above the ex-post delivered capacity in the previous year

Such a rule will reduce administrative burden on both DR providers and CPUC staff, while still retaining oversight abilities in cases where a DR provider underperformed in the previous year or a significant increase in QC is requested.

# CHAPTER 4: Incentive Mechanism

The final component of the proposed process (step 5) is an incentive mechanism that is assessed based on the ex-post delivered capacity relative to committed capacity (which is limited to QC). An incentive mechanism known as the Capacity Shortfall Penalty (CSP) is proposed as an alternative to the current incentive mechanism in the California ISO markets, the Resource Adequacy Availability Incentive Mechanism (RAAIM). Unlike the CSP, RAAIM is assessed on bids relative to a must-offer obligation (MOO). This chapter first addresses RAAIM and the MOO, then introduces the CSP as an alternative.

## RAAIM and the MOO

The RAAIM is assessed based on bids over the course of the AAH, which are indeed the hours in which loss-of-load events are likely to occur. However, the MOO requires resources to bid their net QC in each AAH. This structure generally appears sufficient for traditional dispatchable generation resources that can produce a constant output over many hours; if a natural gas power plant bids 100 MW for five consecutive hours, it is highly likely to deliver that power if called upon to do so.

For this proposal to function as intended, elimination of RAAIM and the fixed MOO is proposed for all DR resources. The California ISO is requested to clarify that DR providers can and should bid their true availability rather than QC value and to ensure that DR providers are not in violation of the ISO tariff for doing so.

DR must be recognized as a variable output resource that must be able to bid according to its actual capability, rather than a fixed MOO. Even if the MOO varies by hour under the slice-of-day framework, it must be able to bid variably across different days with different weather conditions. As such, we propose exemption from the RAAIM. An alternate approach would be to implement a weather-adjusted MOO, which changes with temperature per the ex-ante capability profile developed in step 1. However, the California ISO system is not currently able to implement such a variable MOO.

Simply eliminating RAAIM would retain a static MOO by hour in each month but would eliminate the financial penalties associated with offering a lower value. We recognize that retaining the fixed MOO for DR could put providers out of compliance with the California ISO tariff when offering lower values and suggest a variable MOO be investigated in the future. In the near- to mid-term, eliminating the RAAIM will have the same effect.

Intervals for which a resource has an RA obligation but does not bid are imputed with a BNLI of zero when determining load impacts in step 3 above, providing an incentive to bid.

## Capacity Shortfall Penalty

DR resources are fundamentally different and only some types of DR can deliver sustained constant load impacts over many consecutive hours. Even so, variable DR resources can provide significant capacity contributions. The incentive mechanism differs from the RAAIM by applying to the ex-post measured capacity relative to the committed capacity, defined as the lesser of contracted capacity and QC. In doing so, it accounts for actual performance where applicable through the definition of BNLI. This feature is critical to ensure DR providers cannot avoid penalties under a RAAIM-like system by bidding the contracted capacity value and purchasing the difference in the spot market.

The CSP is defined as the product of any shortfall in demonstrated capacity relative to the contracted capacity, the market price for capacity, and a penalty parameter:

$$CSP = 2\, P \max\left(0.945 Cap_{Com} - Cap_{Dem}, 0\right)$$

where the value of 2 is the parameter that defines the relative intensity of the penalty, P is the price of capacity, 0.945 is the percentage of committed capacity (94.5%) below which the penalty is imposed on demonstrated capacity, $Cap_{Com}$ is the committed capacity, and $Cap_{Dem}$ is demonstrated capacity. The cutoff of 94.5 percent is chosen for consistency with the existing RAAIM structure and mitigates performance risk by allowing slight underperformance that may be attributable to some combination of random conditions and statistical measurement error.

Note the maximum function ensures DR providers face a penalty for delivering below the capacity award, but do not receive a bonus for surpassing it. However, we also propose allowing DR providers to contract for a range of capacity within which they will be paid for their delivered capacity. The DR provider will only face the penalty when the portfolio delivers less than 94.5 percent of the bottom of the range. However, we note that this provision need not be explicitly adopted for the RA program — these can be negotiated in RA contracts with LSEs.

Figure 13 illustrates the effective revenue of a resource relative as a function of its demonstrated capacity, including a provision to be compensated for up to 110 percent of committed capacity for illustration. Below 47.25 percent of committed capacity, the DR provider will owe more in penalties than the contract value; by 0 percent it will pay 94.5 percent of its entire contract value back as a penalty.

**Figure 13. Effective Capacity Revenue under Capacity Shortfall Penalty**

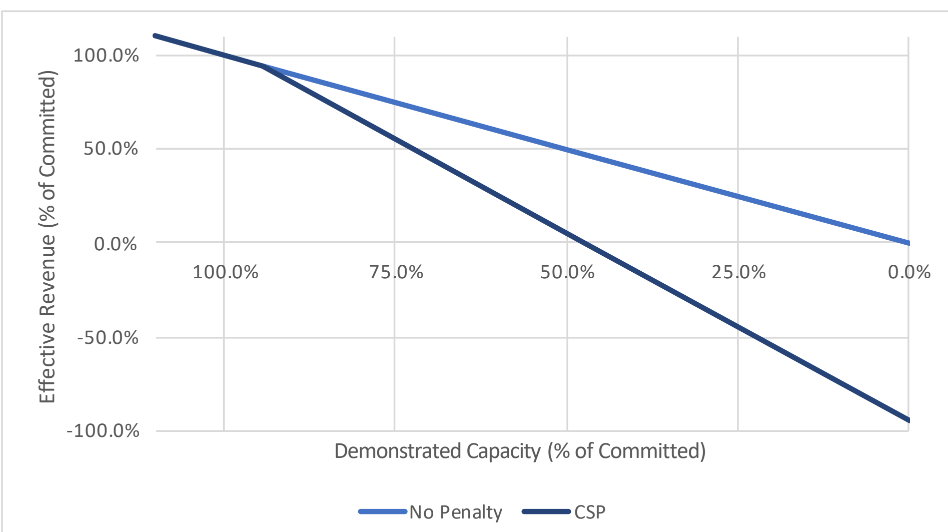

The CSP will apply on a monthly basis to the average capacity value of the resource across the hours (that is, slices) to which it is committed. These include hours with takeback or other negative load impacts. As an example, see Figure 10 and note the difference between the average capacity in the fixed window for the 1-in-2 peak temperature and the variable bidding that better reflects how an economic resource might behave. So long as the average ex-post capacity value (0.47) is no less than the average ex-ante capacity value (0.44), the resource will not face a penalty.

Underperformance risk from unavoidable future uncertainty and randomness can also be actively mitigated through aggregating a DR portfolio as discussed in the following section.

## Capacity Aggregation

Underperformance risk can be mitigated by aggregating delivered DR capacity across a provider's resources before applying the CSP described above. To illustrate, we assume DRPs face the CSP and that DR providers can aggregate their resources that are eligible to provide the same capacity product (for example, system capacity). Consider ten hypothetical DR providers, each with one-hundred resources with 1 MW expected capacity and a standard deviation of 0.4 MW. Each resource is contracted to provide 1 MW of capacity. The total expected value of each DR provider's aggregate capacity is 100 MW with standard deviation 4 MW. A simulation of each provider's aggregate capacity contribution resulted in values from about 94–110, as shown in Table 3.

**Table 3. Simulated capacity deliveries and shortfall with and without aggregation.**

| Delivered Capacity (MW) | Shortfall (No Aggregation) | Shortfall (Aggregated) |
|---:|---:|---:|
| 109.49 | 12.33 | 0.00 |
| 104.96 | 14.03 | 0.00 |
| 102.30 | 15.74 | 0.00 |
| 101.62 | 15.12 | 0.00 |
| 98.84 | 18.31 | 1.16 |
| 97.95 | 15.79 | 2.05 |
| 97.03 | 19.38 | 2.97 |
| 96.06 | 18.06 | 3.94 |
| 95.92 | 19.64 | 4.08 |
| 94.41 | 19.15 | 5.59 |

Without aggregation, faced shortfalls in the range of 10–20 percent, which equates to 20–40 percent of the capacity value of the resources as proposed under the CSP. Notably, all DR providers – including those that overperform in aggregate – face a shortfall because individual resources that underperform are not cancelled out against those that overperform. In contrast, shortfalls are much lower with aggregation, ranging from 0–6 percent.