

Oregon Public Utility Commission

December 2020

Principles of Capacity Valuation

UM 2011 Capacity Investigation



Energy+Environmental Economics



This whitepaper is prepared by:

Ben Shapiro

Zach Ming

Arne Olson

Sumin Wang

Nick Schlag



Contents

- 1. Introduction 1
- 2. A Capacity Valuation Framework..... 1
 - Key Question 1: How much capacity can a resource provide? 1
 - Effective Load Carrying Capability..... 2
 - Heuristics to Approximate ELCC..... 5
 - Summary 8
 - Key Question 2: What is the value of capacity?..... 9
 - Summary 11
 - Capacity Compensation Frameworks 12
 - Summary 15
- 3. Application of the Capacity Valuation Framework 15
 - General Principles and Cross-Cutting Considerations..... 15
 - Capacity Contribution (MW) 15
 - Capacity Value (\$/MW) 16
 - Determining “Peak” Periods 16
 - Renewable Generation 17
 - Storage 19
 - Demand Response 20
 - Hybrid Resources 21
 - Energy Efficiency 22
- 4. Conclusion..... 22



1. Introduction

Capacity is a critical element of an electricity resource portfolio. Capacity reflects the ability to meet demand in all hours, including both peak hours and across a wide range of load and resource availability conditions. This report presents an analytical framework for valuing the capacity contribution of different energy resources as well as considerations on potential frameworks to compensate different resources for the capacity they provide. Energy and Environmental Economics (E3) was engaged by the Oregon Public Utility Commission (OPUC) to conduct this work as considered through UM 2011.

This report outlines the importance of using a consistent set of principles in valuing capacity across all energy resources and use cases to ensure that one technology or customer is not favored over another. However, energy resources have many different characteristics beyond their simple capacity that make a single compensation framework difficult to apply to all energy resources in all cases. To address these topics, the report is structured as follows:

Section 2 provides an overview of two key questions which together create the foundation for determining the value of capacity.

1. How much capacity can a resource provide (MW)?
2. What is the value of capacity (\$/MW)?

Additionally, this section explores potential compensation frameworks that can be implemented that reflect the fundamentals of these two questions.

Section 3 uses this overview of information to explore various applications of capacity value to specific resources and use cases. **Section 4** concludes the report.

2. A Capacity Valuation Framework

Determining the capacity value provided by a resource requires consideration of two distinct questions:

1. How much capacity can a resource provide (MW)?
2. What is the value of capacity (\$/MW)?

These questions are separable and independent. There are different approaches which can be taken to answering them, each with its own benefits and limitations.

A separate but related topic follows— how to compensate a resource for its capacity – that should reflect the answer to the first two questions.

This section of the report discusses each of these topics in turn, describing the primary considerations and distinct approaches relevant to each one.

Key Question 1: How much capacity can a resource provide?

Different resources provide different amounts of capacity relative to their nameplate size, depending upon their operational characteristics, limitations, or other constraints. While no resource can provide “perfect capacity” due to forced outages and maintenance, this theoretical concept of a “perfect capacity” resource



is a useful benchmark for measuring capacity contribution and is widely used across the industry. Comparisons of real-world energy resources to a theoretical perfect capacity resource provide a consistent manner by which to evaluate capacity contributions.

The primary basis to determine the capacity contribution of a resource should be adherence to loss-of-load probability (“LOLP”) principles. This concept is well-established in the electricity sector and has been used for many decades to determine overall capacity needs of a system to comply with a specific reliability standard. For example, many systems are planned to a target loss-of-load expectation (“LOLE”, i.e., the number of days per year with expected loss of load) standard of one day in ten years or 0.1 days per year. If two resources provide the same reliability contribution to the system (i.e. the system yields equivalent LOLE with either resource) then these resources provide equivalent capacity.

Effective Load Carrying Capability

Effective Load Carrying Capability or ELCC is becoming increasingly recognized as the “gold standard” approach to accurately measuring the capacity contribution of a resource using LOLP principles. ELCC is a technology-neutral measurement of equivalent “perfect” capacity of any resource. For example, if solar has an ELCC of 50%, an electricity system with 100 MW of solar would achieve the same reliability as a system with 50 MW of a perfect firm resource.

As shown in Figure 1, ELCC is calculated by 1) calculating system reliability, 2) adding the desired resource to the resource portfolio, and then 3) removing perfect capacity until the original level of reliability is restored. These calculations can be done using any system reliability model that adheres to LOLP principles by calculating reliability over a wide range of system conditions.

Figure 1. ELCC Calculation Process



A resource’s ELCC is equal to the amount of perfect capacity removed from the system in Step 3

As ELCC is calculated using a system reliability model, this measure takes into consideration not only the specific characteristics and capabilities of the resource in question, but also the usefulness of the resource in providing capacity to meet the specific reliability needs of the electricity system in which it is operating.


ELCC Dynamics

Due to complex interactions between non-firm resources such as wind, solar, storage, and demand response, it is difficult to measure the ELCC of an *individual* resource while accounting for these interactions. Some pairings of resources are antagonistic and diminish one another’s capacity contributions. Alternatively, other pairings are synergistic and enhance the resources’ capacity contributions.




Figure 2 below provides an example of an antagonistic interaction, where increasing solar installations on the system shift the net peak later in the day, diminishing the capacity contribution of additional solar resources.

Figure 2. Antagonistic Pairing: Solar + Solar offers diminishing returns to scale



In contrast, Figure 3 depicts a synergistic pairing between solar and storage. In this example, the solar resource alone is able to reduce the net peak by only a small amount given the timing of its production relative to the peak period. The storage resource alone is able to reduce the gross peak further but is limited by the amount of energy it can provide relative to the duration of the peak period. When these resources are combined, however, the resulting reduction in net peak is greater than the sum of their respective individual peak reductions.

Figure 3. Synergistic Pairing: Benefits of Solar + Storage



ELCC Measurements

The ELCC of a resource can be measured in several ways. **Portfolio ELCC** is a measurement of the combined ELCC of all intermittent and energy-limited resources. Alternatively, **First-In ELCC** is a measurement of ELCC as if it were the first and only intermittent or energy-limited resource on the system. This measurement ignores the interactive effects discussed in the previous section. Finally, **Last-In ELCC** is a measurement of the marginal ELCC of a resource after all other intermittent or energy-limited resources have been added




to the system. In contrast to First-In ELCC, Last-In ELCC captures all of the interactive effects with other resources.

The Last-In ELCC of a resource can either be higher or lower than its First-In ELCC. If *higher*, this indicates that there are positive or synergistic interactive effects with the other resources on the system, yielding a “diversity benefit.” Alternatively, if Last-In ELCC is *lower* than First-In ELCC this indicates there are antagonistic interactive effects because the resource has similar characteristics to other resources on the system, yielding a “diversity penalty.”

Figure 4 demonstrates how the individual First-In or Last-In ELCC values of different resources may not sum to the Portfolio ELCC. This is because the Portfolio ELCC incorporates the interactive effects – whether synergistic or antagonistic – of all resources on the system, whereas the First-In ELCC by definition ignores these effects and the Last-In ELCC when measured individually for each may count these effects multiple times.

Figure 4. Comparison of ELCC Measurements



Portfolio ELCC appropriately characterizes the capacity contribution of intermittent and energy-limited resources, which is *important for assessing system reliability*. Alternatively, Last-In ELCC appropriately characterizes the marginal ELCC of the next unit of an intermittent or energy-limited resource. This is *important for procurement decisions* and understanding how new resources will contribute to system capacity needs.

ELCC Data Requirements

Calculating ELCC in LOLP models is computationally intensive and requires a significant quantity of data. From a principled perspective, the model should capture load and resource performance under a wide array of system conditions that could possibly result in loss of load. Because loss of load is relatively infrequent in a well-planned system (i.e. once every ten years), most LOLP models simulate the system for hundreds if not thousands of years using different combinations of load and generation conditions. Most LOLP models

use at least thirty (30) years of historical weather data and eight (8) years of renewable generation profiles in order to accurately converge on statistically significant ELCC values.¹ When historical data is not available, it must be simulated using the best available estimates of its actual productive capability. For renewables, the National Renewable Energy Laboratory (NREL) System Advisor Model is an industry-accepted method for developing renewable generation profiles. For energy-limited resources such as energy storage and demand response, incorporating the duration of these resources and any limitations on how often they can be used/called is critical in calculating ELCC.

Additional ELCC Considerations

It is important to note that the capacity contribution of an energy resource is dependent on more than just a resource's physical capability but also how the resource is operated. To the extent that a resource is dispatchable, such as energy storage, different operational strategies for otherwise identical resources may result in different ELCCs. To the extent that a resource's compensation framework or price signals influence its operation, the two fundamental questions of *"How much capacity can a resource provide?"* and *"What compensation framework should be used to pay for capacity?"* may become inseparable.

It is also important to note that ELCC calculations often require a significant amount of data, particularly for renewable resources in order to estimate their generation patterns over many different years. To the extent that actual production data is not available, estimation techniques must be used, but policymakers should be aware that the modeled ELCC values are only as good as the data used to calculate them.

Heuristics to Approximate ELCC

While ELCC has been increasingly recognized within the industry as the most rigorous and accurate method for measuring the capacity contribution of energy resources, it is also computationally intensive and may not be a practical manner by which to assess capacity for different resources across all use cases. However, several simplified alternatives or heuristics exist that can be used to approximate ELCC. The following sections describe several common heuristics, including their advantages, limitations, and appropriateness for different use cases.

Loss-of-load Probability Heuristic


One common heuristic approach uses hourly Loss-of-load Probability (LOLP) values, which are developed using the same models that produce ELCC values. Hourly LOLP values represent the probability that there will be loss-of-load in a given hour of the year and are based on many simulations of the electricity system under different load and resource conditions. Resources whose generation is coincident with high LOLP hours are estimated to have high ELCC.

Hourly LOLP values are represented as percentage values for each hour of the year, and summing these hourly values across the entire year yields the expected number of hours with lost load in that year. This is often depicted visually as a heat map of the LOLP values for each month-hour (12 months x 24 hours), as shown in Figure 5. The higher LOLP values (darker shades of red) indicate periods with the greatest likelihood of lost load, and therefore the greatest need for capacity.

¹ <https://www.nrel.gov/docs/fy17osti/67501.pdf>



Figure 5. Illustrative LOLP Table



Because hours with high LOLP represent hours in which the system needs additional capacity to improve reliability, calculating the coincidence of a resource’s generation and hourly LOLP values is a reasonable method to approximate ELCC. The calculation process to estimate ELCC using LOLPs is as follows:

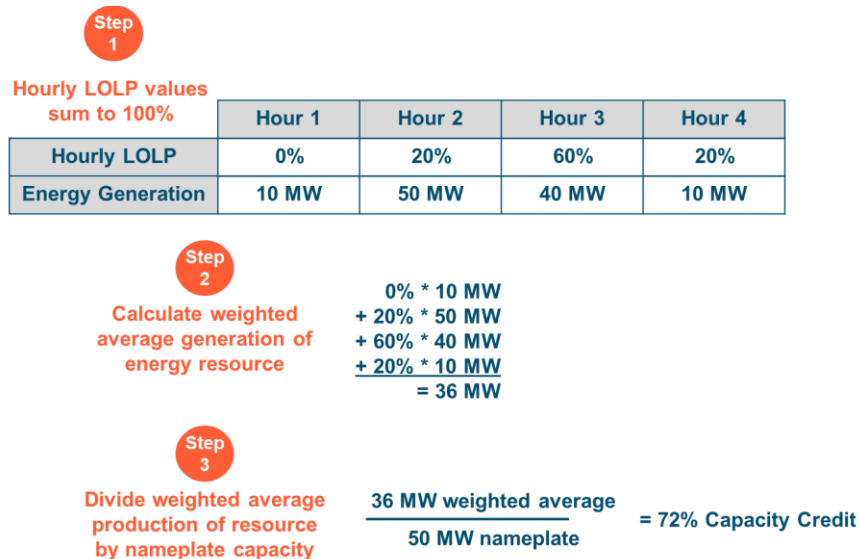
1. Normalize the hourly LOLPs across the year – hourly values should sum to 1.0²
2. Calculate the weighted average production (in MW) of a specified resource over that year, using the normalized LOLP values and the resource’s hourly production profile
3. Divide this weighted average production by the nameplate capacity of the resource to produce a percentage which represents the approximate ELCC % of that resource

A simple mathematical example using only four hours that illustrates this calculation process is provided in Figure 6.

² Normalization of hourly values across the year entails dividing each individual hourly value by the sum of all hourly values in that year such that the values sum to 100%. This produces a set of hourly values that reflect the *relative magnitude* of LOLP in each hour. Energy resources that are available during all LOLP hours are estimated to provide 100% ELCC. Energy resources that are partly available during all LOLP hours are estimated to provide partial ELCC.



Figure 6: LOLP Heuristic Calculation Example



It is important to note that calculating capacity credit using the LOLP heuristic approximates *Last-In ELCC*, as the LOLP should reflect all existing resources that contribute to minimizing the probability of lost load.

Additional LOLP Heuristic Considerations

There are several noteworthy limitations to the hourly LOLP heuristic approximation of ELCC. The weighted average coincidence between hourly LOLP values and energy generation reflects energy generation for all days within a given month at a particular hour, rather than only on the hours where the system actually has loss of load. Days with actual lost load, however, may be correlated with conditions that also produce higher or lower than average resource production. For example, loss-of-load tends to occur on hot days (due to air conditioning), which tend to be sunny, which therefore have high solar generation. Alternatively, loss-of-load may occur on cold days (due to electric heating) which tend to have little to no wind in the Pacific Northwest. The use of production across all hours of the month may fail to fully capture these dynamics.

An effective way to deal with this limitation is through resource class scalars to adjust up or down the capacity value attributed to resources using ELCC. Comparing the ELCC of a class of resources (e.g., solar) to the coincidence of a given resource’s production with LOLP values (as described above) provides a ratio that indicates how well the LOLP approach approximates the ELCC of that resource. Using this ratio (ELCC divided by LOLP coincidence) as a scalar to “true up” the hourly LOLP values is a reasonable manner to more accurately approximate ELCC. For example, if the actual Last-In ELCC of the energy resource illustrated in Figure 6 was 80%, multiplying the hourly LOLP values by 111% (80% / 72%) would allow this heuristic method to more appropriately approximate capacity credit.

Separately, use of the hourly LOLP heuristic as an approximation for ELCC is not as well-suited for energy storage and other energy-limited resources, given that the *length* (duration) of loss-of-load events is not captured directly in hourly LOLP values. For example, while there may be LOLP values across six consecutive hours for a given month (e.g., 4-10pm), this does not necessarily imply that loss-of-load from a *single event* is causing that entire span of LOLP values. The earlier portion of the LOLP in that timespan (e.g., 4-8pm)




could be due to one loss-of-load event, while the later portion of the LOLP (e.g., 6-10pm) could be due to a separate, independent loss-of-load event on a different day. This would mean that a resource capable of providing capacity for four hours, rather than the six-hour period reflected in the LOLP values, would be sufficient to provide 100% ELCC. The use of the hourly LOLP heuristic rather than ELCC directly does not reflect this.

Approximating ELCC using Net Peak Loads

Another approach to approximating ELCC, rather than through coincidence with LOLP hours, is by considering a resource's production during the top X net peak load hours, where X is often in a range of 50-200 hours per year. The use of net peak load hours approximates Last-In ELCC, while the use of gross peak load hours approximates First-In ELCC. Net load is calculated by subtracting all non-dispatchable generation (wind and solar) from gross system hourly loads. The accuracy of this method in approximating ELCC may be limited by factors not captured in this approach, such as changes in the seasonal capability of dispatchable generators to help meet net load.

Figure 7. Net Peak Load Hours (Illustrative)



Summary

As discussed in this section, the primary considerations and recommendations for evaluating the capacity contribution of a given resource are as follows.

- + If two resources each yield an electricity system with equivalent reliability, then these resources provide equivalent capacity.
- + ELCC is the most accurate measure of capacity by measuring this value based on contribution to system reliability.
- + Because ELCC can be computationally complex, approximations of ELCC value can be estimated using heuristic methods, such as the coincidence of a resource's output with hours of high LOLP.
- + Portfolio ELCC appropriately characterizes the capacity contribution of all intermittent and energy-limited resources, which is *important for assessing system reliability*.
- + Last-In ELCC is a measure of the marginal capacity contribution of a resource *after* considering the effects of all other intermittent or energy-limited resources, whereas First-In ELCC is a measure of the marginal contribution *before* considering the effects of these other resources.
- + Last-In ELCC appropriately characterizes the marginal ELCC of the next unit of an intermittent or energy-limited resource, which is *important for procurement decisions* and understanding how new resources will contribute to system capacity needs.

Key Question 2: What is the value of capacity?

The *monetary value* of capacity is distinct and separable from the *quantity* of capacity a resource is able to provide. The primary basis to determine the monetary value of capacity should be adherence to avoided cost principles. These principles are such that a resource should be provided no more compensation than the least cost resource that can be procured by the utility that provides equivalent reliability. In order to establish this monetary value, two questions must be answered:


1. Does the utility need new capacity?
2. How much does capacity cost?

Does the utility need new capacity?

The amount of capacity a utility needs is determined by a specified reliability target, with the most common being a 1-day-in-10-year standard. This is equivalent to a 0.1 days/year loss-of-load expectation (LOLE). Resource planners at utilities ensure this standard is met by adhering to a planning reserve margin (PRM) that provides excess capacity above typical peak loads (1-in-2 median). This allows the utility to maintain reliability in the event of atypical combinations of load and resource availability, such as unplanned forced generator outages, unusually high peak loads (generally due to extremely hot or cold weather), or times when generators must withhold capacity in order to meet operating or contingency reserve requirements. Typical PRMs range from 12-20%, depending on a range of system characteristics.

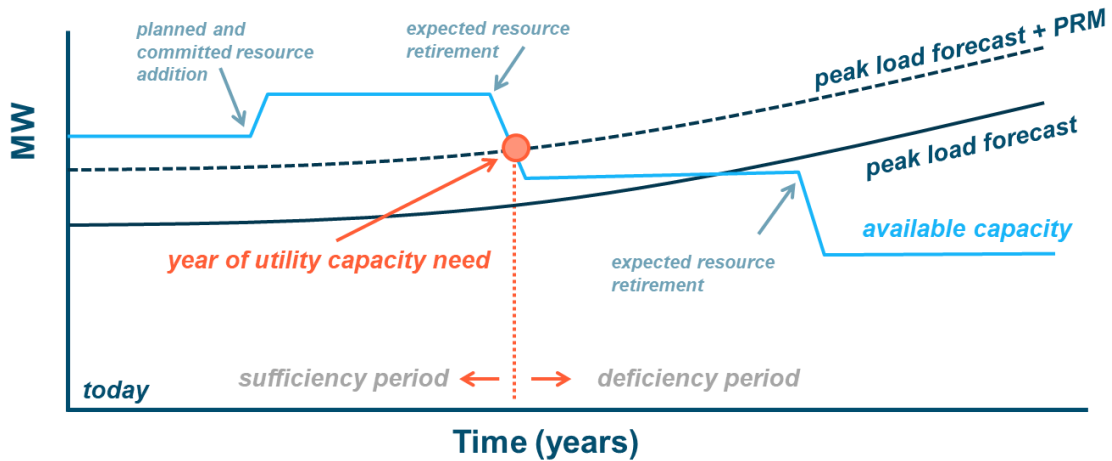
While the PRM establishes the minimum capacity required to meet anticipated needs above peak load estimates, utilities commonly hold reserves in excess of this level due to the additional reliability these reserves provide, the difficulty in accurately predicting peak demand, and the irregular or “lumpy” nature of investments in new capacity.

When planning over multiple years, utilities forecast anticipated load growth as well as anticipated additions and retirements of generation resources (Figure 8). Through this exercise, a deficiency year is identified, in which load plus the required PRM exceeds expected available capacity. Prior to the deficiency year, when capacity exceeds anticipated load plus the PRM, the utility is in a period of capacity sufficiency.



Traditional System

Figure 8. Capacity Availability, Peak Load Forecast, and Planning Reserve Margin Over Time (Illustrative)



When assessing whether the utility is sufficient or deficient in future years, analysis should only consider projects that are fully committed and have no opportunity for deferral or modification. This avoids a reliance on projects that may not actually materialize, which would jeopardize the reliability of the system.

How much does capacity cost?

The cost of capacity is directly linked to whether the utility is in a sufficiency or a deficiency period. Generally, approaches for valuation of capacity in either period use the net resource cost as the starting point. Calculating net resource cost identifies the lowest net cost capacity resource. Net resource cost is equal to the gross cost of the capacity resource less the value of the system benefits it provides and can be compensated for, such as energy and ancillary services, as depicted in the equation below.

Figure 9. Net Resource Cost Calculation

Gross cost of capacity	\$/kW-yr
– System benefits	\$/kW-yr
energy	
ancillary services	
etc.	
= Net cost of capacity	\$/kW-yr

Traditionally, combustion turbines have been the lowest net cost of capacity resource in the electricity system. Other resources should not be used to establish the net resource cost unless they are lower cost or if there are policy limitations on lower-cost resources, such as restrictions on construction of fossil fuel plants.

Capacity Cost in Sufficiency vs. Deficiency Periods

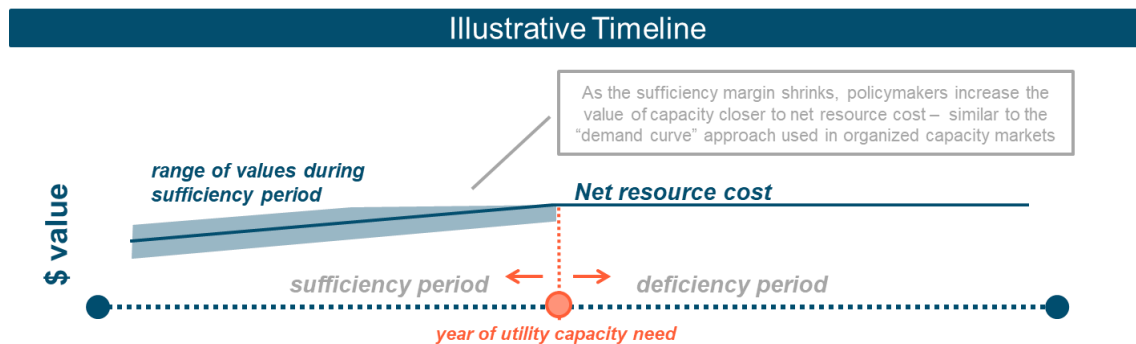
In periods of deficiency, the net resource cost is used directly as the cost of capacity because it reflects the anticipated cost to build new capacity resources. In periods of sufficiency, a common approach to valuing capacity is to use the fixed operations and maintenance cost of the net resource cost resource. This



approach is based on the cost to maintain existing capacity resources such that they are available to ensure system reliability, while also recognizing that the full cost of new capacity resources is an excessive measure of capacity value in times where sufficient resources are available.

Generally, as the sufficiency margin – the difference between available capacity and anticipated load plus PRM – shrinks, policymakers increase the value of capacity closer to full net resource cost, in recognition of the growing need for new capacity resources as the deficiency year approaches. This is similar to the “demand curve” approach used in organized capacity markets.³ Figure 10 provides an illustrative example of capacity value in sufficiency and deficiency periods.

Figure 10. Capacity Value in Sufficiency vs. Deficiency Periods (Illustrative)



Additional Considerations in Establishing Capacity Value

Capacity can be provided either by utilities or by third parties, and it is important to consider equity between these groups in their respective contributions to resource adequacy. One potential inequity is the ability of utilities to earn full cost recovery for capacity procurements *in excess of* the PRM, while third party providers of capacity may be constrained by the PRM. This equity consideration must be balanced against the potential for an economically inefficient outcome where payments are made for capacity in excess of the amount needed to ensure reliability.

Summary

The primary considerations and recommendations for determining the value of capacity provided by a resource are as follows:

- + Determining when a utility will require new capacity is based on load forecasts, the required planning reserve margin, and committed resource additions and retirements.
- + Prior to this capacity need, the utility is in a capacity sufficiency period; after that point, the utility is in capacity deficiency period.
- + The value of capacity in a deficiency period is equal to the net resource cost of the lowest cost capacity resource.

³ Demand curve constructs used in competitive electricity markets (e.g., PJM, NYISO, ISONE) adjust the clearing price of capacity based on how short of long the system is, relative to the reliability standard.

- + The value of capacity in a sufficiency period is equal to the fixed operations and maintenance cost of the lowest cost resource.

Capacity Compensation Frameworks

The framework to compensate energy resources for capacity should ideally reflect both the quantity of capacity (MW) each resource provides and the monetary value (\$/MW) of that capacity. To the extent that the price signals sent by the compensation framework impact how a resource is dispatched, the framework should incentivize resources to dispatch in a manner that maximizes the capacity contribution to the utility system without creating unnecessary requirements.

Because each energy resource is unique along several dimensions (dispatchability, metering, etc.), it is difficult to construct a single compensation framework that is appropriate for all energy resources. In addition, the ultimate compensation framework for each energy resource often involves balancing conflicting factors, including accuracy and tractability. Determining the balance between these factors is a policy judgement.

This report provides an overview of two general capacity compensation frameworks: a “fixed payment” method and a “pay-as-you-go” method, as summarized in Table 1 below. Other compensation frameworks may be reasonable but should be evaluated using the principles outlined in this report.

Table 1. Fixed Payment and Pay-as-you-Go Compensation Frameworks

	Fixed Payment	Pay-as-you-Go
Method	Resources are compensated based on a fixed annual value (\$/yr) that aligns with their capacity credit (MW) and the value of capacity (\$/MW-yr)	Resources are compensated based on production during capacity scarcity hours (e.g., high LOLP hours)
Application	Dispatchable resources where performance can be directly measured. This compensation framework is used in organized capacity markets such as PJM, NYISO, etc.	Non-dispatchable resources such as PURPA renewables and other DER resources
Performance Evaluation	Resources are compensated on a fixed basis as long as they adhere to pre-specified performance requirements	Resources are compensated on a pay-as-you-go basis for production during capacity scarcity hours

Fixed Payments

In a fixed payment compensation structure, resources are compensated based on a fixed annual value (\$/yr) that aligns with their capacity credit (MW) and the value of capacity (\$/MW-yr). In regulated markets, the value of capacity is often established using net resource cost and a sufficiency/deficiency determination (as described in the previous section), while the capacity credit can be determined either through ELCC calculations or a heuristic method such as hourly LOLP-based approximations of ELCC.

An essential component of a fixed payment compensation framework is the inclusion of performance requirements, to ensure that upfront payments for capacity result in actual capacity delivery to the system.



These performance requirements are often implemented such that resources receive financial penalties if they do not perform according to their capabilities.

Fixed payment structures work well for dispatchable resources, where capacity is inherently tied to operational decisions. A dispatchable resource can deliver full capacity to the system by simply being available and delivering energy during the few days a year when capacity is constrained, as opposed to having to generate energy on a daily basis to be fully compensated for its capacity. A fixed payment approach with performance requirements that are dynamically sent to resources during periods of system capacity need can maximize the capacity value of these resources without putting undue burden on the operation of these resources during non-scarcity hours. The primary limitation of this approach is the feasibility of implementation, as it requires the system operator to send a dynamic dispatch signal to a resource and for that resource to respond accordingly.

Pay-as-you-Go

In a Pay-as-you-Go compensation structure, resources are compensated based on energy production during hours of expected capacity scarcity. These hours can either be based on real-time dynamic system need or on pre-determined periods with expectation of high capacity needs (e.g., informed by high LOLP hours). As shown in Table 2 below, there are benefits and drawbacks to each approach.

Table 2. Pay-as-you-Go Implementation Options

	Real-time Dynamic Payments	Pre-determined Time Periods
Method	Compensates resources on a dynamic basis during times of system stress	Compensates resources for performing during pre-determined time periods (e.g., high LOLP hours)
Time Period Frequency	Relatively rare, e.g. only on hot, peak load days	Relatively common, e.g. during all summer afternoon and evening periods
Pros	Properly rewards resources for generating or dispatching when they are needed	Easy to plan for; predictable outcomes
Cons	Does not provide a predictable signal to resources	Rewards resources for performing when capacity it not needed Does not capture the correlation between output of resources on actual peak days ⁴

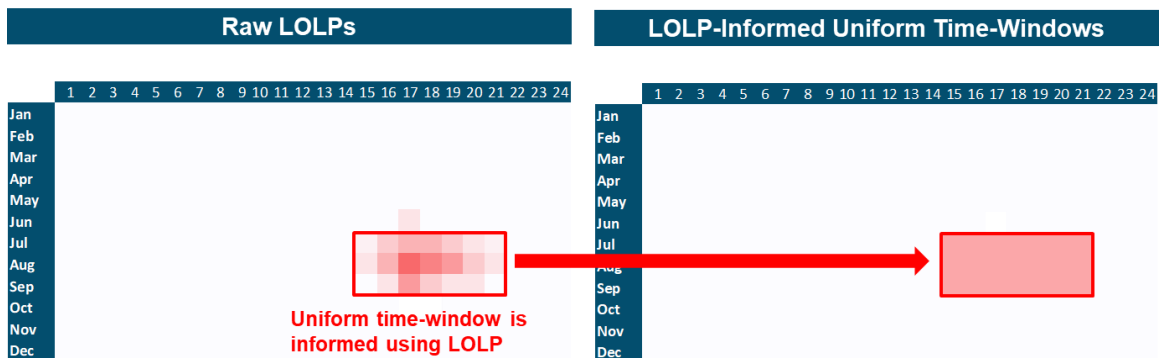
Real-time payments compensate resources dynamically based on when the system is stressed and therefore has the greatest need for capacity. This structure is better aligned with periods of actual (vs. potential) system need, as price signals are communicated in real time. However, this approach is challenging for energy resources to forecast when and how much they will be compensated given the inherent uncertainty in when these periods of actual system need will occur.

⁴ As discussed earlier in this report, one solution to this issue is to implement a “scalar” that adjusts capacity values to account for an ELCC to Resource Output-LOLP-Coincidence Ratio.

Pre-determined time periods compensate resources for their performance during specific time periods, such as high LOLP hours. These periods should be pre-defined for each year of a resource’s economic life (or the contract length, if different) to reflect changes in high LOLP hours over time. This structure is more predictable and therefore easier for resources to forecast but is inherently less aligned with system need given that the pre-determined time periods are based on periods of potential (vs. actual) system need. Additionally, this approach misses important correlations between resource output and conditions that drive system need (e.g., hot days with high peak loads are also generally correlated with high solar production, which is not captured in this approach).

As discussed earlier in this report, hourly LOLP coincidence can be a reasonable heuristic for Last-In ELCC. However, because hourly LOLP values are different for each hour of the day for each month of the year, this can be a complex price signal to send to generators. Alternatively, a common simplification is to consolidate these hourly LOLP values uniformly into a single or multiple peak period(s), thus providing a simpler signal for resources to respond to that still broadly conveys system needs. An illustration of the conversion of hourly LOLP values to a single peak period is shown in Figure 11. There are multiple approaches to determining how to translate LOLP values to a set of peak periods, but the k-means clustering algorithm is one common and reasonable approach.

Figure 11: Illustration of Hourly LOLP Values to Single Peak Period



Contract Length

An additional component of compensation structure is determining how far into the future capacity value should be locked in. Capacity resources are generally capital intensive. In order to incent development of these resources, asset owners often require a degree of certainty on the compensation they will receive over time. This factor is often accounted for in the length of the contract that energy resources sign with the electric utility, as is the case with PURPA. Because the utility tends to be in a period of capacity sufficiency in the near term and capacity deficiency in the longer term when capacity is more valuable, longer contracts are advantageous to third-party energy resources.

Another factor to consider in determining contract length is equity between utility-owned and third-party resources. When a utility constructs a new resource, they are eligible to recover the full costs of that resource over its economic life. To the extent that third-party resources are not offered similar contract terms that cover the full economic life of the resource, this may present an inequity between the two.



Summary

The primary considerations for determining a capacity compensation structure are as follows:

- + Fixed upfront payment structures are suitable for dispatchable resources which can be called upon during the periods of greatest capacity need.
- + Pay-as-you-go structures are suitable for non-dispatchable renewable resources given that their output is weather-dependent.
- + In fixed payment structures, performance requirements may be included to ensure that resources provide the capacity they are being compensated for.
- + In pay-as-you-go structures, payments can be made either based on real-time, *actual* system capacity needs (more accurate) or based on pre-determined time periods of *anticipated* system capacity needs (more predictable); this is a policy tradeoff.
- + Longer contract length can provide both compensation certainty for third-party resources and equity with how utility-owned resources are treated.

3. Application of the Capacity Valuation Framework

This section provides considerations for the application of fundamental capacity concepts to compensation frameworks of specific energy resources or use cases.

General Principles and Cross-Cutting Considerations

Capacity Contribution (MW)

E3 believes the following general principles are reasonable in the determination of capacity compensation structures:

- + Capacity contribution (MW) based on “Last-In ELCC” is consistent with avoided cost principles
 - o Last-In ELCC is consistent with the marginal value that each resource provides and consistent with the utility’s avoided cost of capacity
 - o Last-In ELCC can be calculated using any system reliability model that adheres to LOLP principles by calculating reliability over a wide range of system conditions
 - o All resources on the same electricity system should be modeled using the same ELCC model for consistency
- + For multi-year capacity contracts, locking in capacity contribution (MW) for each future year of the contract provides certainty to third-party resources and equity with utility-owned resources
 - o Determining the capacity contribution (MW) for each year of the contract reflects expected changes to the system such as the decline in Last-In ELCC for resources such as solar, while locking in the values over the life of the contract and providing certainty to the resource owner
 - o To the extent that a third-party resource is eligible to sign a contract equal to the length of the life of the asset, treating contract renewals equivalently to new resource contracts is both equitable with utility-owned resources and does not diminish certainty to third-party resources



- + For dispatchable resources, fixed-payment Last-In ELCC accreditation in conjunction with performance requirements can provide certainty to both the third-party resource in how they will be compensated and certainty to the utility in the capacity value to the system
 - o This approach simultaneously ensures that resources are compensated for performing during critical hours, are not compensated for performing during non-critical hours, and are not required to perform during non-critical hours to earn full compensation of the capacity contribution they provide
- + For non-dispatchable resources, a pay-as-you-go compensation structure based on hourly LOLP values adjusted by the ratio of Last-In ELCC to LOLP coincidence can appropriately compensate third-party resources without placing undue performance requirements on these resources
 - o This approach ensures that resources are appropriately compensated for their capacity without placing undue performance requirements on these resources given that it is difficult to determine if non-performance during critical hours is due to a factor such as weather that is already accounted for in Last-In ELCC
 - o Locking in these LOLP periods upfront for each year of the contract (including a forecast of how they change over time to reflect resource additions, retirements, and changes in load) can provide certainty to third-party resources while adhering to avoided cost principles

Capacity Value (\$/MW)

E3 believes the following general principles apply across all energy resources and use cases in determining the monetary value of capacity:

- + The monetary value of capacity value (\$/MW) based on the net cost of capacity during periods of capacity deficiency is consistent with avoided cost principles
- + The monetary value of capacity value (\$/MW) based on the annual fixed operations and maintenance cost of the lowest net cost resource during periods of sufficiency is consistent with avoided cost principles. As the system approaches deficiency, increasing the value toward the full net cost of capacity may be consistent with the incremental reliability provided by these resources. Ensuring that capacity is accounted for explicitly and separately from other avoided cost attributes can avoid double-counting and over compensation to third-party resources.

Determining “Peak” Periods

Several of the compensation structures considered in this section are based upon “peak” periods defined by periods with high LOLP. Historically, peak periods have been closely if not entirely aligned with peak demand, but the increasing penetration of renewable energy is increasingly de-linking these two. Instead, peak demand is now much more closely tied to periods of “net” peak demand, where net peak is defined as demand minus renewable generation. As discussed in Section 2, it is often simpler to consolidate “raw” LOLP values into peak periods – rather than using the raw values directly – in order to provide a more straightforward price signal for resource owners. Whether and how to consolidate LOLP values into these periods is a policy decision that must balance accuracy and tractability.

To consolidate raw LOLP values into peak periods, one reasonable approach is to use the k-means clustering method. The question of how many peak periods or “clusters” to set is a policy decision. More periods will relay more accuracy in capacity pricing to resources but at the expense of additional complexity. For the



remainder of the discussion on compensation structures, E3 uses LOLP values directly (rather than peak periods). However, this discussion could be amended to provide uniform capacity value for production in all hours of the pre-determined peak period if this simplification is deemed a useful policy decision to enable more straightforward implementation of these compensation structures.

Renewable Generation

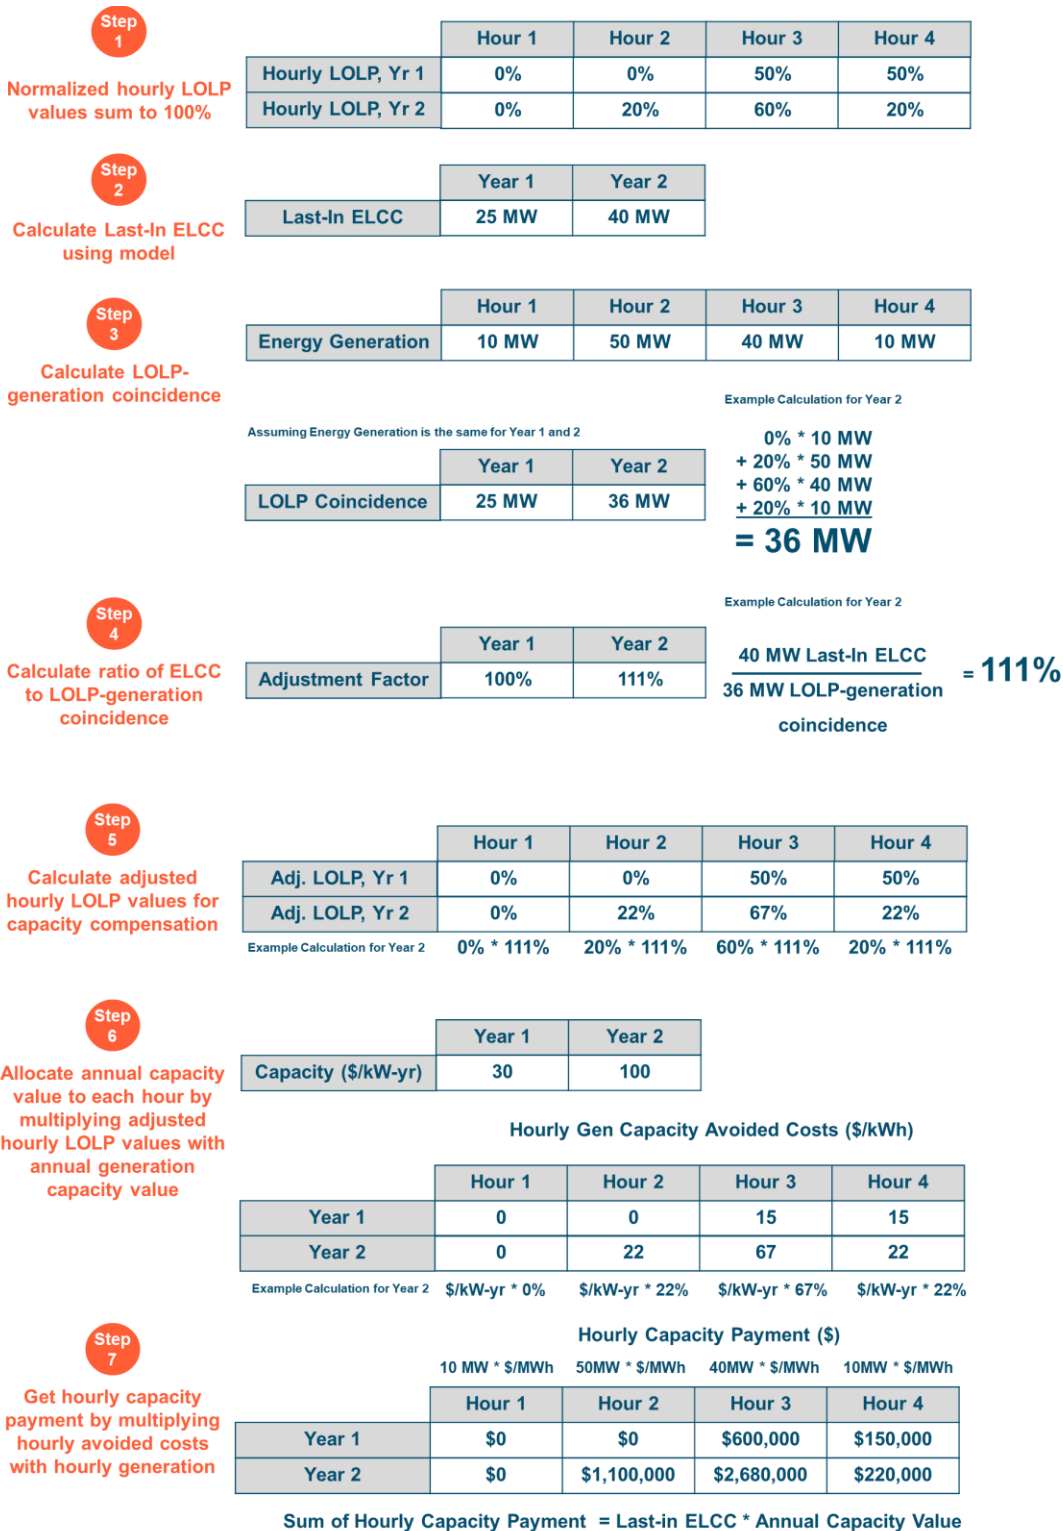
Compensating non-dispatchable renewable via a pay-as-you-go structure can provide appropriate compensation that balances accuracy and simplicity. Non-dispatchable renewable generation includes solar, wind, run-of-river hydro, and any other resource that is not directly controllable by the resource owner due to weather-dependency. This capacity compensation framework can be applied to either utility-scale or behind-the-meter resources. The \$/MWh hourly payment values can be set equal to adjusted hourly LOLP values multiplied by the monetary value of capacity. The adjustment to the hourly LOLP values can be based on the ratio of Last-In ELCC to LOLP-generation coincidence. The end result of this compensation structure is that a resource that generates as expected will be compensated equivalently to its Last-In ELCC.

Renewable Generation	
Capacity contribution	Last-In ELCC, attributed via a pay-as-you-go compensation structure
Compensation framework	Pay-as-you-go compensation structure with hourly compensation values set proportionally to normalized hourly LOLP values, adjusted by the ratio of Last-In ELCC to hourly LOLP-generation coincidence

Figure 12 provides an example numerical calculation of the process for adjusting hourly LOLP values by the ELCC to LOLP-generation coincidence ratio. While this simple example uses four hours, the actual calculation would use upward of 288 values to represent each month-hour (12x24) combination. The adjusted hourly LOLP values should be multiplied by the monetary value of capacity (\$/MW) to determine hourly compensation in the pay-as-you-go compensation framework.



Figure 12: Numerical Example of Hourly LOLP Adjustment Process for Pay-as-you-Go Compensation Structure



Compensating renewable generation resources through a pay-as-you-go structure can be a compelling option because unlike dispatchable resources, the generation of non-dispatchable renewable resources is dictated solely by weather conditions (i.e., the wind and the sun). It is therefore more difficult to compensate these resources through fixed payments with performance requirements when a lack of performance is difficult to attribute to weather variability that is already captured in the ELCC value or some other reason. A pay-as-you-go structure provides an appropriate incentive for owners of these resources to ensure their generators are well-maintained and available to generate when weather conditions permit.

The payments to each generator would be calculated by multiplying the annual monetary value of capacity (\$/MW) by the adjusted hourly LOLP value, as illustrated in the example calculation shown in Figure 12. Using this approach, each renewable resource would receive an ultimate annual payment that totals to the product of their Last-In ELCC and the monetary value of capacity (\$/MW). This adjustment process inherently captures factors such as correlations between renewable production and high LOLP hours that may not be captured in the unadjusted hourly LOLP heuristic approach.

The Last-In ELCC value that could be used for each resource would reflect the specific characteristics of that renewable resource (e.g., fixed axis vs. tracking solar generators) and could be calculated for each year of the resource's life. Calculating the ELCC value upfront for each year of the contract requires projected LOLP values in future years to accurately reflect the expected capacity contribution over time. The utility must make these projections available for 10 – 20 years (or longer) depending on contract length.

While calculating Last-In ELCC values for individual renewable projects can be a computationally intensive exercise, one simplification is to use a “library” approach where the electric utility calculates Last-In ELCC value for a variety of resources with different characteristics and then assigns each actual renewable resource to a resource in the library with similar characteristics. Defining representative resource classes should capture a meaningful distinct set of characteristics such as plant design, age, and geography for renewable resources and duration and efficiency for energy storage. A resource class could be as small as three (wind, solar, and four-hour storage) or could encompass tens of representative resources.

Including more resources in the library is more accurate but comes at the tradeoff of additional complexity. E3 believes that sufficient modeling should be performed such that the Last-In ELCC assigned from the library to each individual renewable resource is within 5% of its true Last-In ELCC.

The compensation structure that results from this potential approach is a \$/MWh capacity payment which resources receive for generation during the times of greatest system capacity need. These periods are based on the highest LOLP hours for each future year of the resource lifetime and ultimately compensate each renewable resource based on its Last-In ELCC (MW) multiplied by the monetary value of capacity (\$/MW).

Storage

Compensating storage resources through fixed annual payments with performance requirements can provide appropriate compensation that balances accuracy, efficiency, and fairness. The fixed payment can be based on the product of the storage resource's Last-In ELCC and the monetary value of capacity.



Storage	
Capacity contribution	Last-In ELCC, attributed via a fixed payment compensation structure
Compensation framework	Annual fixed payment (\$/MW) with performance requirements

Unlike renewable resources, storage assets are dispatchable and primarily provide capacity to the system based on operator decisions as opposed to weather conditions. Further, requiring a storage resource to dispatch (cycle) each day in order to earn its full capacity payment – as would be required under a pay-as-you-go construct – puts an undue burden on storage given that it incurs costs from cycling even though this provides no capacity value on many days that do not face a capacity need. Furthermore, a pay-as-you-go structure has the potential to significantly compensate resources even if the resources do not generate on days of capacity need.

A fixed-payment compensation structure for storage resources would be based on two factors:

- + The capacity contribution as measured by the resource’s Last-In ELCC (MW), and
- + The monetary value of capacity (\$/MW)

For storage resources entering into a long-term contract with the electric utility, providing both the Last-In ELCC and monetary value of capacity at the outset of the contract for each year (i.e. the values can change on a year-to-year basis) contributes to certainty for the developers of third-party resources. The Last-In ELCC in each year of the contract life should be based on expectations of capacity additions and retirements on the utility system, with the storage resource in question attributed the Last-In ELCC value in each year after considering the impact of other resources to reducing LOLP.

As discussed earlier in this report, fixed payment structures paired with performance requirements help to ensure that compensation is commensurate with the value these resources provide to the system. Giving the utility the ability to dispatch the storage resource based on its capabilities (e.g. taking into account duration limitations of the storage) using a day-ahead signal, helps to realize the storage resource’s full capacity potential. In practicality, the utility would only need to allow the utility to dispatch storage approximately ~30 days per year.⁵ If storage resources do not operate when dispatched by the utility, a financial penalty would be assessed, requiring the resource owner to return a portion of the fixed payment.

Demand Response

Compensating demand response resources through fixed annual payments with performance requirements can provide appropriate compensation that balances accuracy, efficiency, and fairness. The fixed payment would be based on the product of the demand response resource’s Last-In ELCC and the monetary value of capacity. The performance requirements would be based on the inherent capabilities of the demand response resource and would be the identical limitations used in calculating its ELCC. This compensation framework is very similar to the potential framework for energy storage.

⁵ The 30-day figure may be modified but is intended to reflect approximately the number of hours per year used for different capacity valuation practices across North America (e.g., 200 hours/year).



Demand Response	
Capacity contribution	Last-In ELCC, attributed via a fixed payment compensation structure
Compensation framework	Annual fixed payment (\$/MW) with dispatch requirements commensurate with the limitations assumed in the Last-In ELCC calculation

The capacity value of demand response is highly dependent on how these resources are dispatched, similar to storage. However, demand response resources often have limitations on how many times per year a utility can “call” on a customer to reduce load or for how many hours a customer will reduce load when called upon. In this sense, demand response is more limited than energy storage. In order to maximize the capacity value of demand response, the electric utility must save demand response calls for the times of greatest capacity scarcity.

A fixed price compensation framework based on Last-In ELCC and the monetary value of capacity compensates demand response for the marginal capacity it provides to the system. As with storage, penalties for DR resources’ non-performance when called help ensure that fixed payments are not made without actual capacity being provided when needed. Importantly, the assumptions of call or dispatch limitations should be consistent between assumptions used in the calculation of Last-In ELCC and the performance requirements it is bound by. These include both the frequency with which the resource can be called and the duration of each call.

Hybrid Resources

Hybrid resources share characteristics of two distinct individual resources: renewables and storage. This presents the option of compensating such resources for their capacity contribution based on the generating resource (i.e., renewable portion), or separately compensating the components based on their individual characteristics. The decision as to which compensation framework is more appropriate can be made either by the resource owner or by the utility.

Hybrid Resources	
Capacity contribution	Last-In ELCC, attributed solely via a pay-as-you-go compensation structure <i>or</i> in conjunction with a fixed payment compensation structure
Compensation framework	Two options: <ol style="list-style-type: none"> 1. Pay-as-you-go compensation structure for combined system 2. Pay-as-you-go structure for the renewable portion of the system and a fixed payment structure for the storage portion of the system

If electing to compensate the hybrid resource based solely on the pay-as-you-go structure, this would be implemented identically as with renewable generators. The potential downside of this compensation structure is that it creates a large burden on storage to cycle on a daily basis which may not be necessary to realize full capacity value.

Alternatively, if electing to compensate the resource separately based on the renewable and storage components separately, the renewable portion would receive compensation through a pay-as-you-go construct, while the storage portion would receive a fixed payment in conjunction with performance



dispatch requirements, identical to the structure described above for standalone storage resources. It is important that the generation associated with the fixed payment be subtracted from the generation being compensated via pay-as-you-go in order to avoid double counting capacity compensation. Note that compensating hybrid resources *entirely* on a fixed price basis is likely not a practical option given that the generation portion of the resource remains dependent on variable weather conditions, and therefore developing performance requirements is very difficult as discussed in the renewable generation section.

Energy Efficiency

Unlike many of the resources previously discussed, energy efficiency is often not directly meterable. Instead, its capacity value is likely best calculated through models and assumptions about its performance. The capacity contribution of energy efficiency can be based on the Last-In ELCC of the resource. Further, capacity valuation of energy efficiency is often only used in cost-benefit analysis of energy efficiency incentives and programs and is not used in actual compensation of these resources.

Energy Efficiency	
Capacity contribution	Last-In ELCC
Compensation framework	Value of energy efficiency for cost-benefit analysis purposes could be based on the net present value of the product of a) the forecasted Last-In ELCC by year over the life of the measure and b) the monetary value of capacity for each year of the resource’s life

As with the other resources described above, the capacity contribution of energy efficiency resources should be determined using an industry standard LOLP model to calculate Last-In ELCC. This calculation can be done upfront to establish the ELCC value of the energy efficiency resource for each year of the life of the resource.

Given that energy efficiency is not typically metered, a valuation based on expected (assumed) performance is a reasonable approach for these resources. The value of energy efficiency resources can therefore be determined by taking the net present value of the product of a) the forecasted Last-In ELCC, and b) the forecasted annual capacity values. This approach considers the expected Last-In ELCC contribution of the energy efficiency resource as well as the expected value of capacity in each year, including expected resource sufficiency and deficiency.

4. Conclusion

Properly evaluating the capacity contribution of a given resource and compensating the resource owner for that contribution requires consideration of a number of issues. The capacity valuation framework presented in this report – centered around the questions of capacity contribution (MW), value of capacity (\$/MW), and suitable compensation structure – attempts to contemplate these considerations in a consistent fashion across all energy resources and use cases.

As with many policy decisions, designing the appropriate capacity compensation structure for different resources requires some level of balance or compromise between economic efficiency, equity between resources, transparency and tractability, and ease of implementation. The framework presented in this



report provides a set of principles to apply across technologies to develop reasonable compensation structures for the capacity provided.

While the resulting considerations for compensation structure vary by resource type based on operational characteristics and limitations, all are unified around the general principles described within the valuation framework. These principles primarily include: assigning resources a capacity credit or contribution (in MW) based on their marginal or “Last-In” effective load carrying capability in each year of the resource or contract life; the importance of providing certainty around that contribution for the duration of the contract by locking in values upfront; determining the monetary value of capacity based on the net cost of either the least cost new capacity resource or that resource’s operations and maintenance costs, depending on whether the utility does or does not need new capacity; and ensuring that resource operations used to develop the compensation structure are consistent with the assumptions used to evaluate the capacity contribution of that resource.

