Purpose

The purpose of this document is to provide an explanation of the changes that PJM has proposed to the reserve market design and Operating Reserve Demand Curves (ORDCs) at the Energy Price Formation Senior Task Force (EPFSTF). For the purposes of this document, the term "reserves" will be used to refer to reserves used in real-time system operations, not those associated with the system planning process.

PJM's Goals for Reserve Market Enhancements

Below are several high-level goals that guide PJM's proposal. While there are many details that have been and will be discussed within the stakeholder process, PJM's proposals to date and those that will come in the future will align with these principles:

1. Reserve and energy prices reflect system conditions and appropriately value scarcity.
2. ORDCs reflect the reliability value of reserves.
3. The actual reserve capability on the system is accurately determined.
4. Resources assigned reserves will provide them when deployed.
5. Market power is mitigated.
6. Social welfare is maximized.[1]

Today, these goals are not always met. For example, energy and reserve prices do not always align with system conditions. This can occur for a number of reasons, including the ORDCs in use today not accurately reflecting the reliability value of reserves and the difficulty in determining the amount of reliable reserves on the system due to the voluntary nature of some synchronized reserves. It can also be argued that under today's design, social welfare is not maximized because the ORDCs do not reflect the true value of reserves. PJM seeks to enhance the reserve market design by ensuring that its proposal will achieve these goals.

The Relationship Between the Capacity Market and Operating Reserve Markets

In the capacity market, PJM procures capacity on behalf of load to ensure that there are adequate resources available during the delivery year to meet energy and reserve needs based on forecasted system conditions. Capacity market revenues are intended to ensure that the remaining fixed costs of a resource are covered, given an expectation of what the Energy and Ancillary Services Market revenues will be. Capacity market revenues are not intended to cover the real-time operating costs of a resource, such as startup, fuel or variable maintenance costs. Real-time operating costs are compensated through the energy market via the locational marginal price (LMP) and in some cases, uplift payments. The LMP reflects the marginal cost of serving load.

Throughout the day-ahead and real-time operating day, the reserve markets procure sufficient generation capacity so that PJM can operate reliably given the inherent uncertainty on the system. This uncertainty includes but is not limited to uncertainty related to resource performance and load forecast error. The reserve market clearing price is

[1] Maximizing social welfare is the objective function of the market clearing algorithms. The goal of this objective function is to optimally allocate resources for energy and reserves such that the final allocation simultaneously maximizes the benefit to consumers and the revenues to suppliers. This is done by maximizing the difference between the consumer's willingness to pay for a product and the bid production cost of cleared supply.

intended to reflect the marginal cost of maintaining those reserves in the same way the LMP reflects the marginal cost of producing energy. This cost is separate and distinct from a capacity market payment. The primary marginal cost associated with providing reserves is lost opportunity cost (LOC). LOC is the profit that a resource foregoes in the energy market by not generating to provide reserve capacity, despite energy market prices incentivizing injection onto the system. Capacity market revenues are not intended to cover these costs.

By capturing the LOC in the reserve market clearing price, reserve markets also provide incentives for resources to lower their output and provide reserves instead of providing energy. Absent a payment for reserves, resources would not have the incentive to provide them because they could maximize revenues by providing energy instead.

Proposed Enhancement: *Consolidation of Tier 1 and Tier 2 Reserve Products*

This proposed change is motivated by the need to enhance the accuracy of PJM's reserve measurements and the reliability of response in addition to reducing the significant out-of-market payments that exist today in the reserve markets. Measuring reserves accurately has always been a PJM focus, but it received dedicated attention following an event in September 2013.[2] During this event, PJM calculated that it had 1,655 MW of synchronized reserves on the system; however, when those reserves were requested to deploy, the response within the first 10 minutes was approximately 200 MW. After an hour of spinning reserve deployment, only 400 MW of resources responded to their reserve assignment. This event, in addition to other similar but less severe events, has heightened PJM's focus on its reserve calculations, and in response to these events, PJM made several enhancements focused on improving its estimation of Tier 1 synchronized reserves.[3] These improvements include the following:

1. Deselecting asset classes that are physically unable to provide reserves. These include resources such as wind, solar, nuclear, etc. Based on submitted data, resources in these classes were receiving Tier 1 estimates but were physically unable to meet their estimates.

2. Incorporating the Degree of Generator Performance (DGP) in reserve calculations. The DGP statistic is a variable that PJM calculates based on how well a generator is following its energy basepoint. It was previously only used in the energy dispatch to estimate how closely a unit would follow a basepoint being sent. Using this statistic in reserve calculations resulted in more conservative Tier 1 estimates for resources that are not following their energy basepoint.

3. Continued tracking and monthly reporting of Tier 1 estimates versus actual event performance. These statistics are provided at each Operating Committee meeting.

Despite PJM's efforts to improve the accuracy of its reserve calculations, performance of the Tier 1 product relative to estimates is not good. In the 2017 State of the Market Report, and as provided in Figure 1, the PJM independent

---

[2] Section 4.1, http://www.pjm.com/~/media/library/reports-notices/weather-related/20131223-technical-analysis-of-operational-events-and-market-impacts-during-the-september-2013-heat-wave.ashx.

[3] PJM classifies all synchronized reserve capability as either Tier 1 or Tier 2 reserves. Tier 1 MW are the incidental 10-minute reserves present on the PJM system due to economic dispatch. Tier 1 is provided by a resource that is on-line, following economic dispatch, and capable of increasing its output beyond its desired energy MW following a call for a synchronized reserve event. Tier 2 is provided by resources that are synchronized to the grid and operating at a point that deviates from economic dispatch (including condensing mode) to provide additional synchronized reserve not available from Tier 1 resources within 10 minutes.

market monitor (IMM) shows an average response of 60.1 percent for Tier 1 with a low end of 14.3 percent. The 60.1 percent reflects the response rate based on the adjusted estimates using the DGP parameter, not the actual estimate based on offer data. The response rate based on unadjusted estimates would be even lower.

Figure 1 Synchronized Reserve Events 10 Minutes or Longer, RTO Reserve Zone: 2017 [4]

Spin Event (Day, Time)	Duration (Minutes)	Tier 1 Estimate (MW Adj by DGP)	Tier 1 Response (MW)	Tier 2 Scheduled (MW)	Tier 2 Response (MW)	Tier 2 Penalty (MW)	Tier 1 Response Percent	Tier 2 Response Percent
Mar 23, 2017 06:48	24	926.8	549.6	742.8	559.1	183.7	59.3%	75.3%
Apr 8, 2017 11:53	10	1,222.6	827.2	879.3	828.7	50.6	67.7%	94.2%
May 8, 2017 04:18	10	1,325.6	976.3	335.1	298.5	36.6	73.6%	89.1%
Jun 8, 2017 03:39	10	974.4	726.7	575.7	522.4	53.3	74.6%	90.7%
Sep 4, 2017 20:03	15	476.3	68.1	601.0	563.8	37.2	14.3%	93.8%
Sep 21, 2017 14:15	16	305.8	217.4	1,253.9	1,037.3	216.6	71.1%	82.7%
2017 Average	14.2	871.9	560.9	731.3	635.0	96.3	60.1%	87.6%

Under today's design, the Tier 1 product is voluntary and only compensated upon response to an event. This leads to an inconsistent response rate that makes it impossible to estimate accurately. As a result, PJM system operators bias Tier 1 estimates due to a lack of confidence that the reserves will be provided when requested. As shown in the preceding table, this reaction is rational given the low response to Tier 1 relative to what is estimated. As long as this level of uncertainty exists, it is difficult to reduce the operator's intervention into the reserve calculations and still maintain reliability.

Tier 2 reserve commitments show much better accuracy and reliability with an average response rate of 87.6 percent. These resources are compensated at the clearing price regardless of whether there is an event or not and face penalties for nonperformance. These two design elements result in much stronger performance that is needed for operational certainty.

PJM believes that consolidating the Tier 1 and Tier 2 reserve product into one uniform synchronized reserve product will solve multiple issues that exist in the reserve markets today. This unified product will:

1. Be assigned based on the market solution that maximizes social welfare
2. Be obligated to respond based on the assigned quantity
3. Be compensated at the applicable clearing price for the assigned megawatt amount
4. Face the existing penalty if the resource does not respond during an event

By applying these standards across all synchronized reserve resources, PJM expects the following benefits:

1. More accurate reserve calculations that require less operator intervention
2. More reliable reserve assignments that will improve synchronized reserve performance
3. Consistent compensation for all resources providing the same service
4. More accurate energy and reserve pricing due to improved synchronized reserve measurement

[4] IMM 2017 State of the Market Report, Section 10: Ancillary Services, Table 10-20 Synchronized reserve events 10 minutes or longer, tier 2 response compliance, RTO Reserve Zone:2017.
http://monitoringanalytics.com/reports/PJM_State_of_the_Market/2017/2017-som-pjm-sec10.pdf

Proposed Enhancement: *Locational Reserve Assignments and Nodal Reserve Pricing*

Under today's market rules, PJM models the Mid-Atlantic and Dominion (MAD) Reserve Zone. This reserve zone has its own requirement that can be met by either reserves within the zone itself or with reserves that can be imported from the rest of PJM. The ability to import reserves into the area is determined by the most-limiting interface between the rest of PJM and the MAD Reserve Zone. For example, if the MAD requirement is 1,000 MW and there is only 300 MW of available flow on the limiting interface, 700 MW of reserves will need to be committed within the MAD Reserve Zone to ensure that there are 1,000 MW of reserves deliverable to it. Other ISO/RTOs use similar reserve zone modeling.

Aside from operating reserve adequacy, the most important consideration in the commitment of reserves is to ensure that when they are deployed, they do not cause overloads on any critical transmission facilities or limiting interfaces.

PJM currently addresses this concern by modeling the MAD Reserve Zone, but that model has some shortcomings:

1. Assuming the limiting interface is defined correctly, the current model solves the problem too conservatively.
2. Modeling multiple reserve zones is difficult, especially when there is overlap between reserve zones.
3. A truly dynamic model that recognizes the impact of reserve deployment on all transmission constraints cannot be implemented using this technique, because reserve zones need to be pre-defined in the technical software before they can be implemented.

As stated previously, it is critical that reserves are assigned to resources that, once deployed, do not create overloads on the transmission system. As long as all of the limiting transfer interfaces and applicable reserve zones can be defined in advance and implemented simultaneously, the current reserve zone modeling method is an effective, yet conservative, way to ensure that the assigned reserves do not create overloads once deployed.

However, defining all of the relevant transmission constraints and reserve zones beforehand is impossible, as unforeseen occurrences such as a transmission element or generator trip may create the need for a transmission constraint or reserve zone. Therefore, the current model is often too rigid. Even if all needed reserve zones could be modeled in advance, the model is still too conservative as it does not utilize the full transmission model to model flows into the reserve zone. This generally has the effect of limiting the amount of reserves that can be imported into the zone in order to satisfy the reserve requirement, which increases the amount of reserves that must be located within the reserve zone.

Today, the individual resource distribution factors are not used to determine the impact of reserves outside of the MAD Reserve Zone on the most-limiting interface. Instead, an average distribution factor model plus a small margin is used due to the need to predefine, on a per megawatt of reserve basis, the impact of reserve outside of the MAD Reserve Zone on the most limiting interface. While this is a step forward from the initial implementation where it was assumed that 100 percent of the deployed reserves outside of the reserve zone flowed across the limiting interface, it falls short of adhering to the actual transmission model and is inherently suboptimal.

PJM is currently investigating a more sophisticated model for locational reserves. This model differs from the current model in the following ways:

1. Categories of transmission facilities and limits are identified in advance for reserve modeling as opposed to a specific facility or monitored element contingency pair.

2. For those identified facilities, the assigned reserves are treated as deployed energy when transmission constraints are controlled. This ensures that the Security Constrained Economic Dispatch (SCED) application will control constraints such that the deployment of reserves will not create overloads. This process will not affect power balance.
3. The deployed reserves, like energy, will impact each identified transmission facility based on the distribution factor of the assigned generator. This allows for a more optimal use of the transmission network.
4. Only one reserve demand curve will be used per product. There is no need for an explicit reserve zone demand curve because there is no subregion of PJM that has a specific reserve requirement.
5. Reserve prices will be nodal like energy prices and will be calculated consistent with the network model.

If a model such as this is implementable, PJM will propose such a model in the future. PJM believes this model provides the following benefits:

1. It directly addresses the actual operational problem regarding the locational assignment of reserves rather than the approximate model that is used today.
2. It is a significant step forward in terms of modeling the dynamic changes to the reserve needs across the system as they occur in real time.
3. More complete utilization of the transmission system will result in lower costs to consumers as compared to the existing model, all other things being equal.
4. Nodal reserve prices provide more accurate price signals for the value of reserves in a specific area.
5. It removes administrative complexity in the market by removing the need for specific reserve zones in favor of a simpler, more direct, model.

Should this model not be implementable, PJM will propose to expand the existing static reserve zone model so that there are more pre-defined reserve zones that can be swapped in or out on a more frequent basis. PJM believes that the optimal frequency would be as frequent as operationally needed, but given the administrative complexity of this model, daily changes to the reserve zone are more realistic.

Proposed Enhancement: *Implement Operating Reserve Demand Curves*

Today PJM uses penalty factors to determine the maximum willingness to pay for reserves and ultimately set energy prices during shortage conditions. The current demand curve for synchronized reserve is shown in Figure 2. The synchronized reserve megawatts demanded in the red portion of the curve, labeled Step 1, are determined by the real-time megawatt output of the single largest online contingency. This quantity criterion is the minimum reserve requirement and has been in place and remained unchanged since 2012, when shortage pricing was implemented. The price of Step 1, $850/MWh, is based on analysis of the out-of-market make-whole payments made for reserves from an operating event in 2007.

The blue portions of the synchronized reserve demand curve, Steps 2A and 2B, were both added more recently (in 2017 and 2014, respectively). Step 2A was implemented in response to FERC's order that PJM implement transient shortage pricing. The purpose was to add a smaller step on the curve to avoid system volatility due to large swings in price for small changes in reserve amounts that would have occurred with just Step 1. Step 2B was added as a result of a package that was approved by PJM members that originated in the Energy and Reserve Pricing and Interchange Volatility Senior Task Force (ERPIV). The simple purpose of this optional step was to create the ability to extend the reserve requirement when PJM operators took actions to schedule additional reserves during conservative operations. To date, this step has not been used.

Figure 2 Current Synchronized Reserve Demand Curve

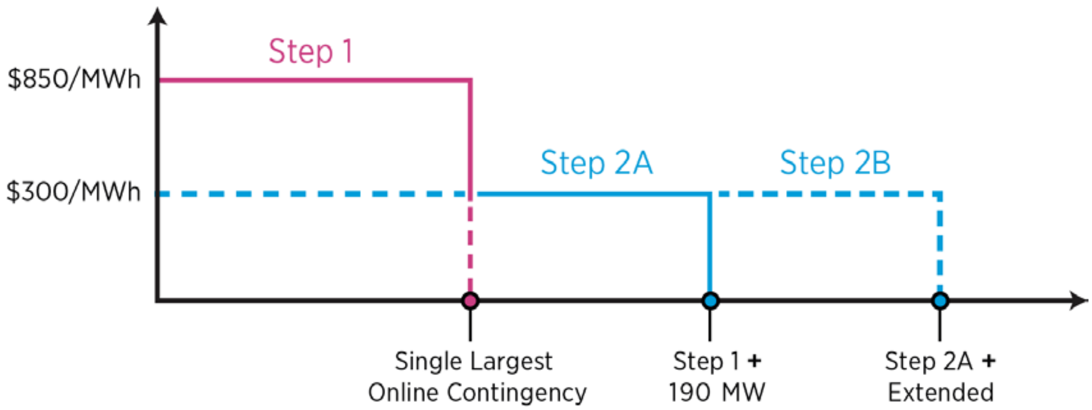

Steps 2A and 2B are both priced at $300/MWh. This number is based on analysis of the offers of quick-start resources at the time that Step 2B was added in 2014.

PJM believes that these penalty factors have the following drawbacks:

1. The curves do not value reserves highly enough.

   In operations today, PJM will take actions in excess of $850/MWh to maintain reserves. Actions such as calling on generators with costs above $850/MWh, deploying pre-emergency and emergency demand response with offers in excess of $1,800/MWh, and a voltage reduction action, which has the effect of reducing load, are taken to preserve the reserves up to the current requirement. Any time PJM takes an action to maintain reserves at a cost above this level, reserve prices do not reflect the cost of reserves.

2. Reserves beyond the identified requirements are not valued.

   Reserves in excess of the currently defined requirements increase system reliability because they allow the system to quickly respond to additional uncertainty beyond the loss of the largest unit, such as load forecast error, renewable generation forecast errors, and so on. This value is not compensated today, and PJM believes it should be recognized.

To address these drawbacks, PJM has proposed replacing the current synchronized reserve demand curve with the ORDC. The ORDC incorporates the minimum reserve requirement in existence today but extends the current curve to include reserves greater than the minimum reserve requirement such that additional reserves are both scheduled and valued when it is rational to do so. Figure 3 is an example curve proposed by PJM.

Figure 3 Example of a Proposed Operational Reserve Demand Curve

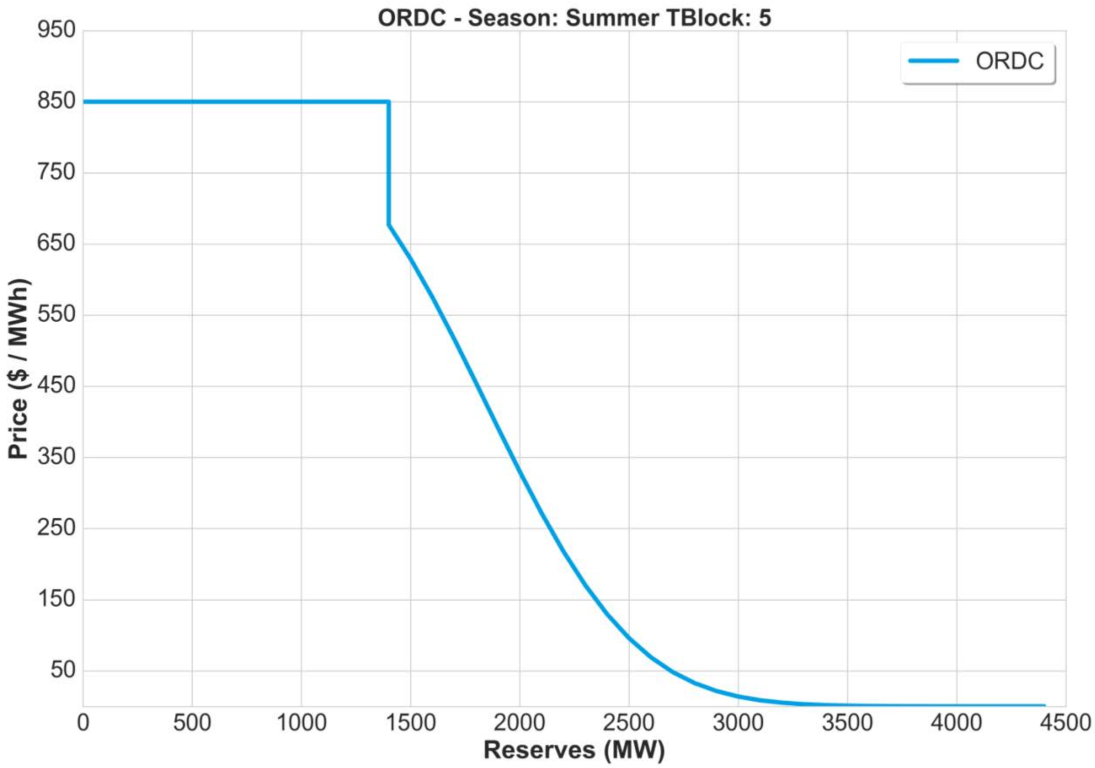

The tail on this demand curve reflects the additional value of reserves beyond the minimum reserve requirement. The minimum reserve requirement is valued at the maximum price for reserves, or $850/MWh. The value of reserves beyond the minimum reserve requirement is determined by multiplying the maximum price on the curve of $850/MWh and the probability of not meeting the minimum reserve requirement (PBMRR) in real-time given the procurement of a specific amount of reserves on a forward basis. The PBMRR function is derived based on historic observed forecast errors for load, solar and wind generation, and thermal generator performance.

Under normal conditions, if there was no uncertainty on the system, PJM would be able to schedule reserves up to the minimum requirement without the risk of falling short of that requirement. This is the assumption that is used in scheduling reserves today. The additional reserves or ramping capability of the system beyond the minimum requirement is assumed to be valueless. However, there are various types of uncertainty on the system that can result in the need to use the additional reserves in order to ensure that the minimum reserve requirement is met.

For example, under today's rules, assume PJM schedules supply to meet forecasted load and its minimum reserve requirement of 1,475 MW of synchronized reserves. If the load forecast is perfect, PJM's minimum reserve requirement is met without issue. However, if the actual load is 100 MW above the forecast or higher, some of the 1,475 MW of synchronized reserve that was procured must be utilized to cover the additional load, which can lead to the inability to meet the minimum reserve requirement.

Based on Figure 3, the probability of not meeting the minimum reserve requirement of 1,475 MW when 1,575 MW are assigned is approximately 70 percent. This can be interpreted as the probability that the combined forecasts error for load, solar and wind, and thermal generator performance will be 100 MW or higher is about 70 percent. If the maximum willingness to pay to meet the minimum reserve requirement is $850/MWh, then it follows that the willingness to pay for 100 MW of additional reserves, or 1,575 MW, to prevent a shortage of the minimum reserve requirement would be $850/MWh * 70 percent, or approximately $600/MWh.

It is important to note that the curves that PJM has presented to date do not address the maximum value of reserves. Discussion on this topic was prioritized for later in the stakeholder process. PJM believes that the maximum level of the ORDC should, at a minimum, ensure that all reserves available to system operators can be captured in the clearing price. This translates to a maximum ORDC level of at least $2,000/MWh. This would ensure that when the system is unconstrained, and generation with offers up to $2,000/MWh are dispatched and set price, the cost of actions taken to maintain reserves are fully captured in the market clearing prices. In addition to increasing the maximum level of the ORDC, PJM believes that adopting the proposed methodology to construct ORDCs for synchronized and primary reserves will have the following benefits:

1. Reserves in excess of the minimum reserve requirement will be appropriately valued based on their benefit to system reliability.
2. PJM will assign additional reserves, when economic to do so, that will result in fewer instances when the minimum reserve requirement is not met.

Conclusion

PJM continues to believe that accurate reflection of operator actions in energy and reserve market prices is a critical component to the continued success of the competitive electricity markets. The proposals described in this paper will enhance the consistency with which prices reflect the value of the resources being counted upon to maintain the energy and reserve requirements on the system. While there are further details required to constitute a complete proposal, PJM believes those explained within this summary are the most critical to ensuring the reserve markets work efficiently. PJM looks forward to working with the stakeholder community through the established process to further refine these details over the coming months. To that end, PJM welcomes stakeholder feedback on the concepts and details described in this paper as the process proceeds.