Operating Reserve Demand Curves (ORDCs)

Emily Barrett

Sr. Lead Market Design Specialist

RCSTF
August 13, 2025

Reserve Design

PJM is building a proposal around a suite of day-ahead and real-time reserve products, each of which will need an ORDC to reflect the value of maintaining these reserves for system operations.

DAY AHEAD	REAL TIME
Day Ahead Reserves	
30 Min. Reserves – Updated	**30 Min. Reserves – Updated**
10 Min Ramp/Uncertainty Reserves – Up and Down Reserves	**10 Min Ramp/Uncertainty Reserves – Up and Down Reserves**
Primary Reserves (PR)	**Primary Reserves (PR)**
Synchronized Reserves (SR)	**Synchronized Reserves (SR)**

Day Ahead Reserves
Addresses uncertainty DA that does not need to be carried into RT.
Inclusive of the need to have energy and reserve commitments to meet the next day load forecast (energy gap)

30 Min. Reserves – Updated
Replacement reserves for contingency events and capturing uncertainty

10 Min Ramp/Uncertainty Reserves – Up and Down Reserves
Addresses net-load ramping needs and allows PJM to meet interval by interval energy and reserve needs

Reserve products would not be nested

Primary Reserves (PR) (10-minute reserves)– 150% SR
Contingency reserves and ACE recovery

Synchronized Reserves (SR) (10-minute reserves) – 100% MSSC*
Contingency reserves and ACE recovery

SR/PR would remain nested

RECALL

30 Min. Reserves – Updated
Replacement reserves for contingency events and capturing uncertainty

10 Min Ramp/Uncertainty Reserves – Up and Down Reserves
Addresses net load ramping needs and allows PJM to meet interval by interval energy and reserve needs

Reserve products would not be nested

Primary Reserves (PR) (10-minute reserves)– 150% SR
Contingency reserves and ACE recovery

Synchronized Reserves (SR) (10-minute reserves) – 100% MSSC*
Contingency reserves and ACE recovery

SR/PR would remain nested

An ORDC is an administrative curve that defines the relationship between the quantity of reserves and the maximum price at a given reserve level.

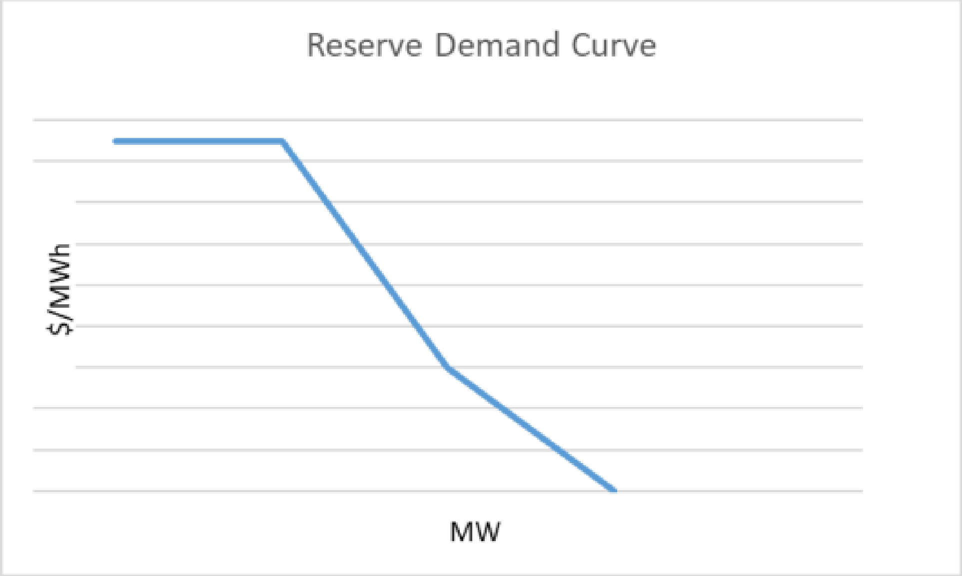

- Y-axis is the reserve penalty factor in $/MWh that sets the maximum price for reserves
- X-axis is the reserve quantity in MW
- The willingness to pay increases as the system approaches reserve shortage or as reserve shortages become more severe. This is sometimes called “Scarcity Pricing.”

Operating Reserve Demand Curves Design Principles

To be effective ORDCs should:

- Align market requirements with reliability requirements.
- Accurately and transparently quantify reserve needs within the market.
- Be informed by operational risks and the costs to mitigate such risks.
- Ensure all cost-effective measures are exhausted before shedding firm load.
- Reflect the value of reserve services as a function of operational conditions.

FERC Statement on Energy Price Formation for Reserve Deficiencies

"To the extent that actions taken to avoid reserve deficiencies are not priced appropriately or not priced in a manner consistent with the prices set during a reserve deficiency, the price signals sent when the system is tight will not incent appropriate short- and long-term actions by resources and load."

-- *Federal Energy Regulatory Commission on [Energy Price Formation (ferc.gov)]*

PJM's Existing Operating Reserve Demand Curves

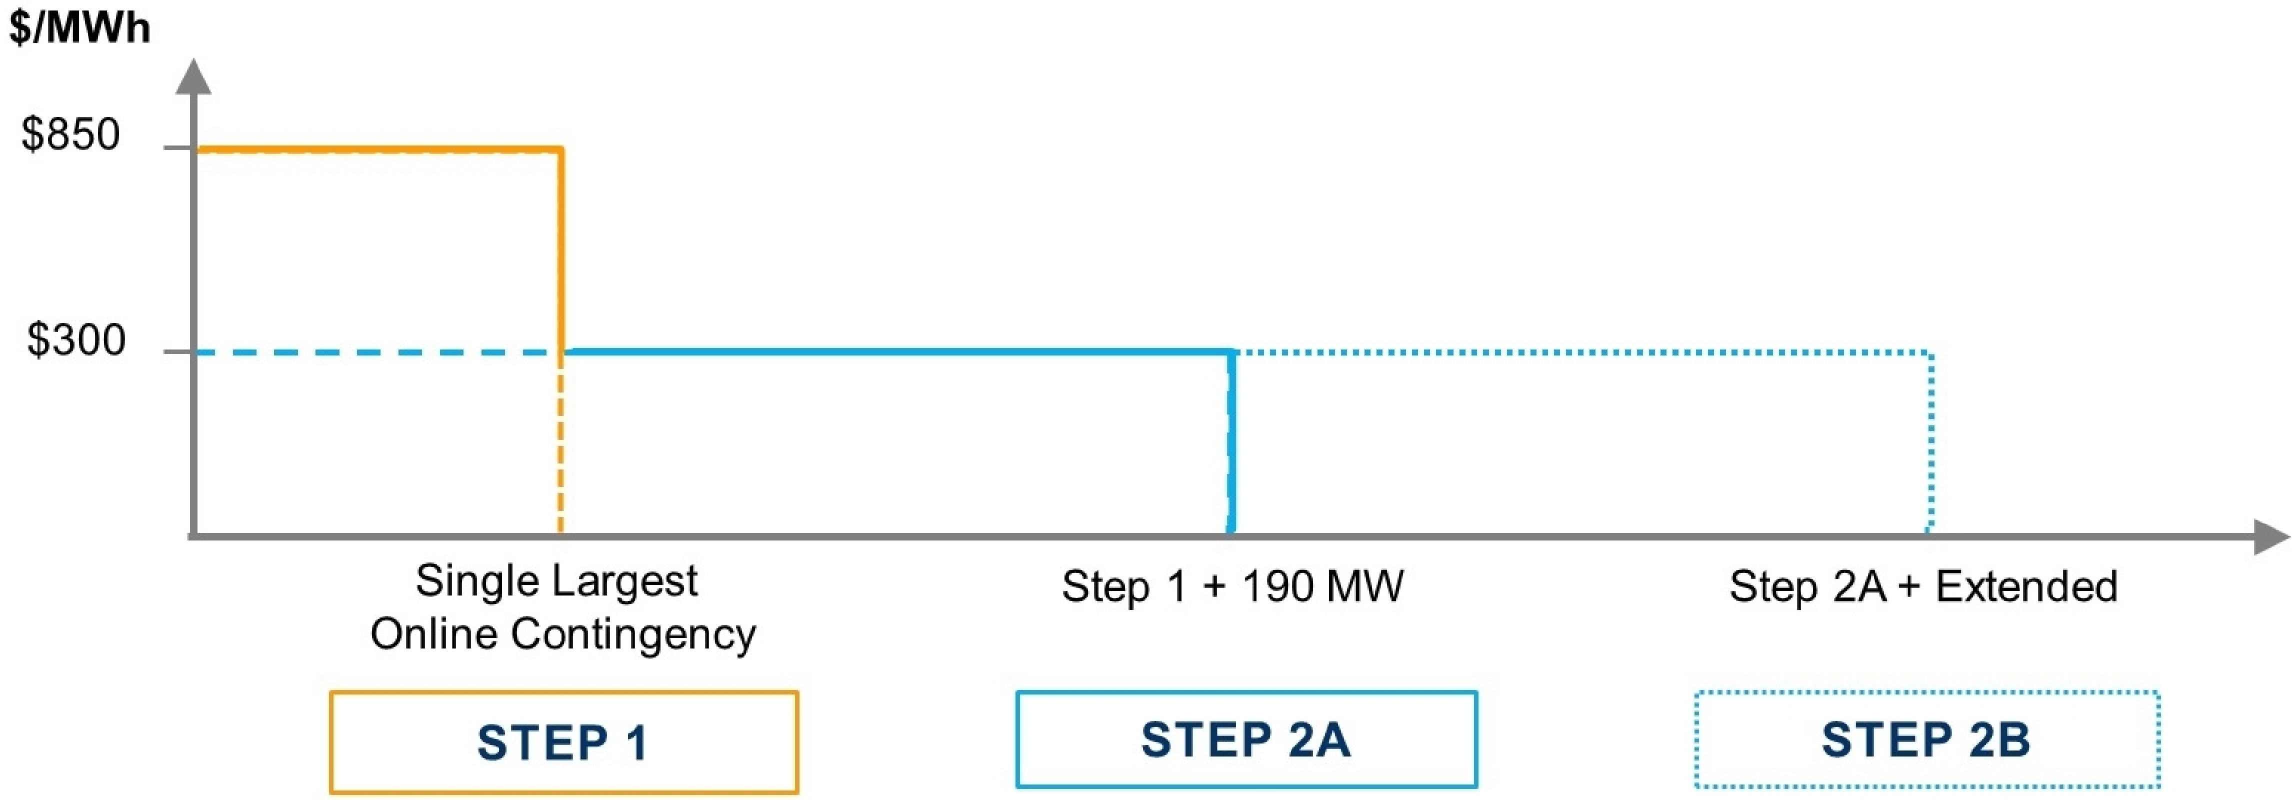

PJM's Concerns About the Existing ORDCs

Stale Operational Data	PJM's existing ORDC penalty factors are based on lost opportunity cost information from an event in August 2007 and do not accurately reflect current operational reality.
Market Inefficiency	As we consider reforms to PJM's reserve markets, our ORDCs need to be updated to reflect the reliability value being provided by each service, to ensure a cohesive market design and to promote market efficiency.
Lack of Alignment	The existing ORDCs do not align with operational needs and the reliability actions that need to be taken to avoid loss of load.

How are we thinking about the ORDCS?

Reserve Quantity (X-Axis)	Reserve Penalty Factor (Y-Axis)	ORDC Shape
• Reserve quantities are dictated by the drivers that lead to the need for reserves, such as forecast uncertainty or managing the risk of unit loss. • Reserve quantities for each reserve product should be based on the reserve needs that the service addresses.	• The reserve penalty factor is driven by the value of reserves – or the maximum price the market will pay for reserves – at each quantity. • The value of reserves may be predicated on system conditions.	• The ORDC's shape will affect pricing dynamics. • Highly vertical ORDC steps can lead to abrupt swings in prices. • The ORDC shape can also be used to represent uncertainties in the system and the trade offs between cost and reliability risk.

Synchronized and Primary Reserve Demand Curves

- In considering reforms to PJM's ORDCs, we should begin with the Synchronized and Primary Demand Curves.
- For clearing reserves in the co-optimization, particularly in PJM's Real-Time Energy Market, SR and PR will have the highest penalty factor(s).
- This provides an anchor for how the other ORDCs fit with the SR and PR Demand Curves.

SR/PR MW Quantities

SR and PR are cleared to manage the risk of losing a large generation unit and align with PJM's reliability obligations as defined by NERC. The factors that drive the procurement of SR and PR are not changing.

- The SR and PR Reliability Requirements are set as a function of the largest contingency to meet our NERC reliability obligations.
- The PR Reliability Requirement is set as 150% of the SR Reliability Requirement.
- The SR and PR Extended Reserve Requirements (i.e., Step 2 of the ORDC) may be reconsidered with the introduction of new reserve products.

SR/PR Penalty Factors

SR and PR are cleared in a co-optimization with energy by market clearing engines which consider the trade-off between going short reserves at the cost of the penalty factor or incurring the cost of clearing reserves, which includes reserve offers, lost opportunity costs and – in the case of unit commitment algorithms – commitment costs.

- To maintain reliability on the system and to avoid reserve shortage, PJM operators routinely commit additional generation resources.
- To make these costs transparent and to promote market efficiency, the penalty for going short reserves in PJM’s markets should be consistent with these costs.
- PJM is therefore exploring using historical resource commitment costs to inform SR and PR penalty factors.

Actions to Remediate Reserve Shortage and Load Shed

The purpose of clearing reserves through the markets is to position the system to make those reserves available. This positioning can entail costs, including:

1. **Lost opportunity costs**, or the costs incurred by resources to provide reserves in lieu of providing energy.
2. **Commitment costs**, or the costs incurred to commit resources, such as start up and no load costs, to make reserves available.

An effective ORDC should ensure that 1) all more cost-effective actions are taken prior to engaging more expensive emergency procedures up to and including Manual Load Dump and 2) that those costs are made transparent to the market. ORDC penalty factors must therefore be high enough to trigger these remediation actions rather than allowing us to go short reserves unnecessarily.

Progressive actions taken during a Capacity Emergency*

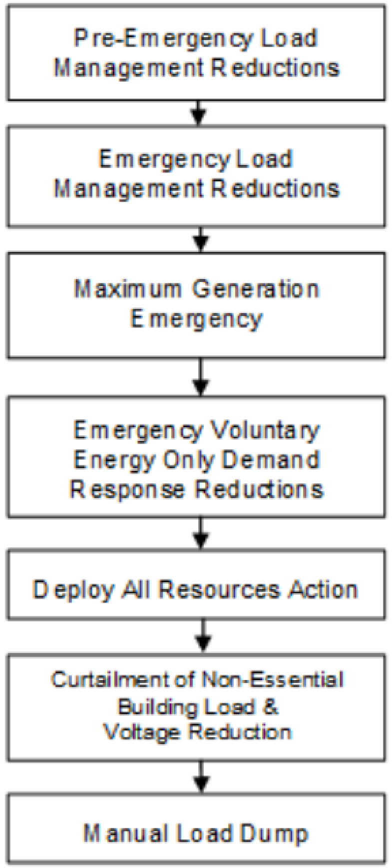

Figure taken from PJM Manual 13: Emergency Operations Section 2.3 Capacity Emergencies

Using Commitment Costs to Inform Reserve Demand Curves

Why consider using commitment costs to inform reserve penalty factor?

- Committing resources that can be started in real-time is the primary operational means of mitigating reserve shortage.
- PJM's market clearing engines use reserve penalty factors to evaluate the relative cost of recommending resource commitments and going into reserve shortage.
- Using commitment costs to anchor reserve prices would help to align operational actions with market outcomes and ensure that available cost-effective measures can be taken before entering reserve shortage.

Using Commitment Costs to Inform Reserve Demand Curves

How could commitment costs be calculated?

- Commitment costs would be based on the sum of the start-up and no-load costs for resources that meet some minimum flexibility criteria (e.g., fast start capable, capable of being committed intraday by IT-SCED) over a defined historical operational period (e.g., the last delivery year).
- These costs would be evaluated for available resources for every hour of the historical period based on the schedule the resource was operating on.
- A distribution would then be calculated of these costs.
- To date, we have started looking at the interquartile range of fast start capable resource commitment costs as an option for anchoring different points on the reserve demand curves.

PJM's Existing SR/PR Demand Curve Shape

- PJM's existing ORDCs value all reserve shortages below the minimum Reliability Requirement at the same penalty factor.
- Whether we are 1 MW or 100 MW short of the requirement the penalty factor is the same, even though the reliability situation is different.

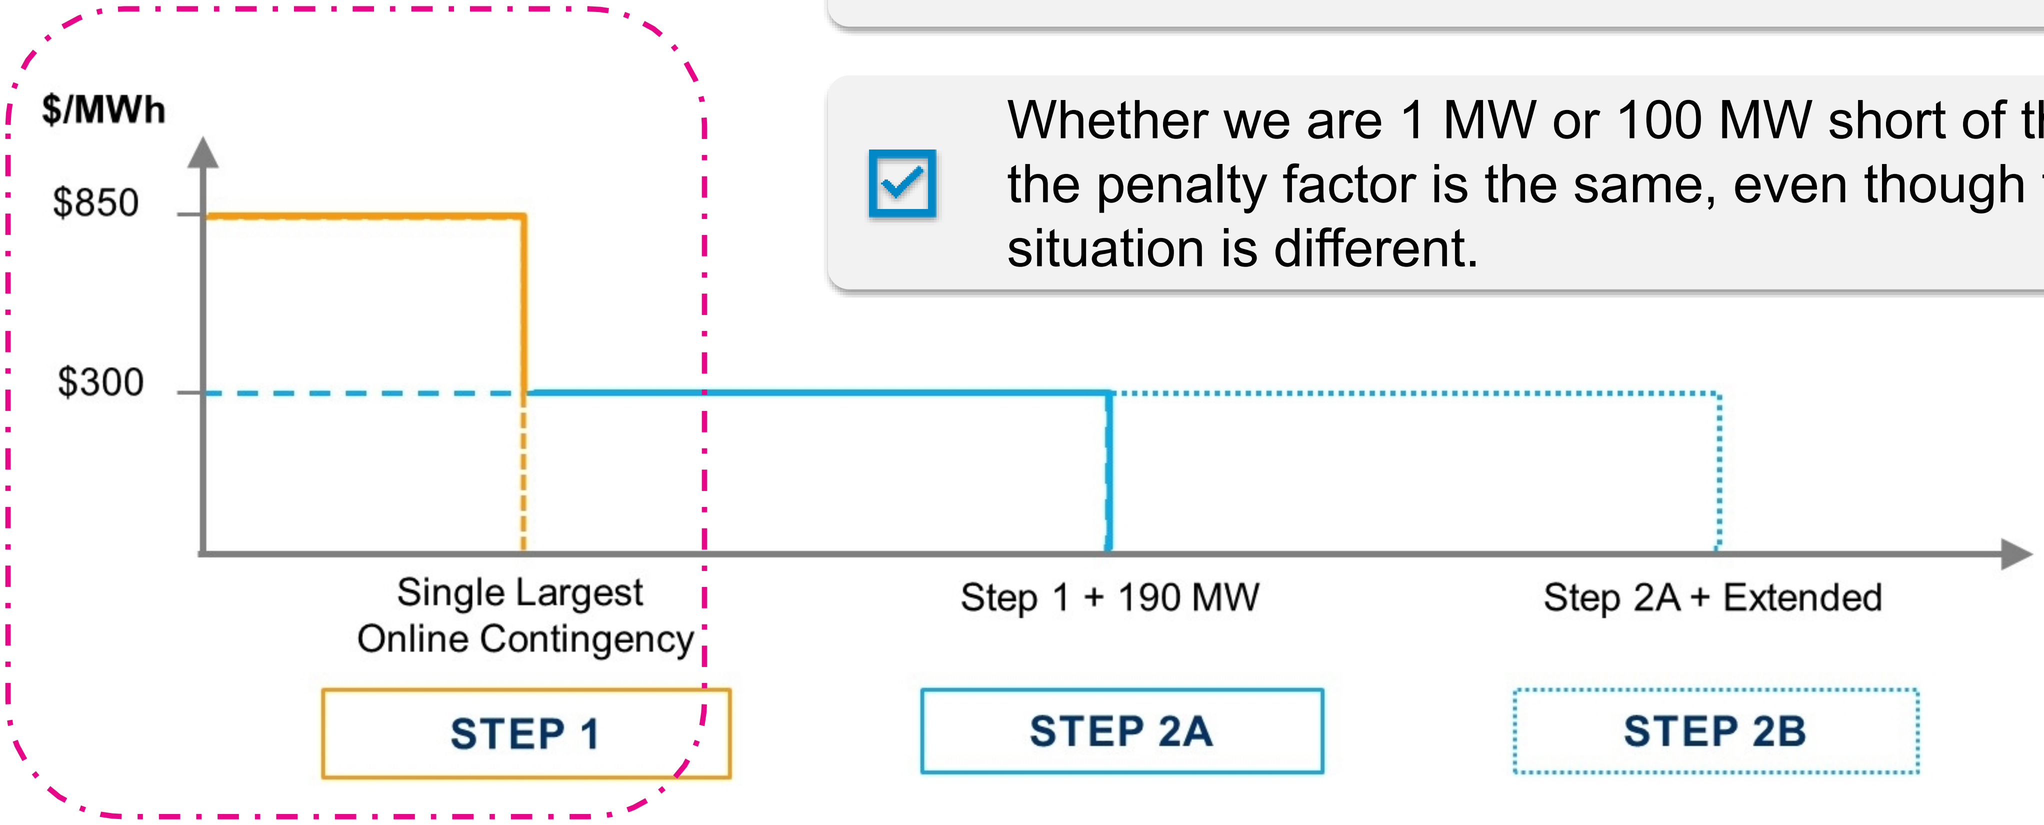

PJM's Existing SR/PR Demand Curve Shape

- Above the minimum Reliability Requirement, there is a single additional step in the ORDC, but it was not designed to manage uncertainties or the probability of going short SR/PR.
- These aspects will be a part of the 10-Min Ramp/Uncertainty Reserve design.

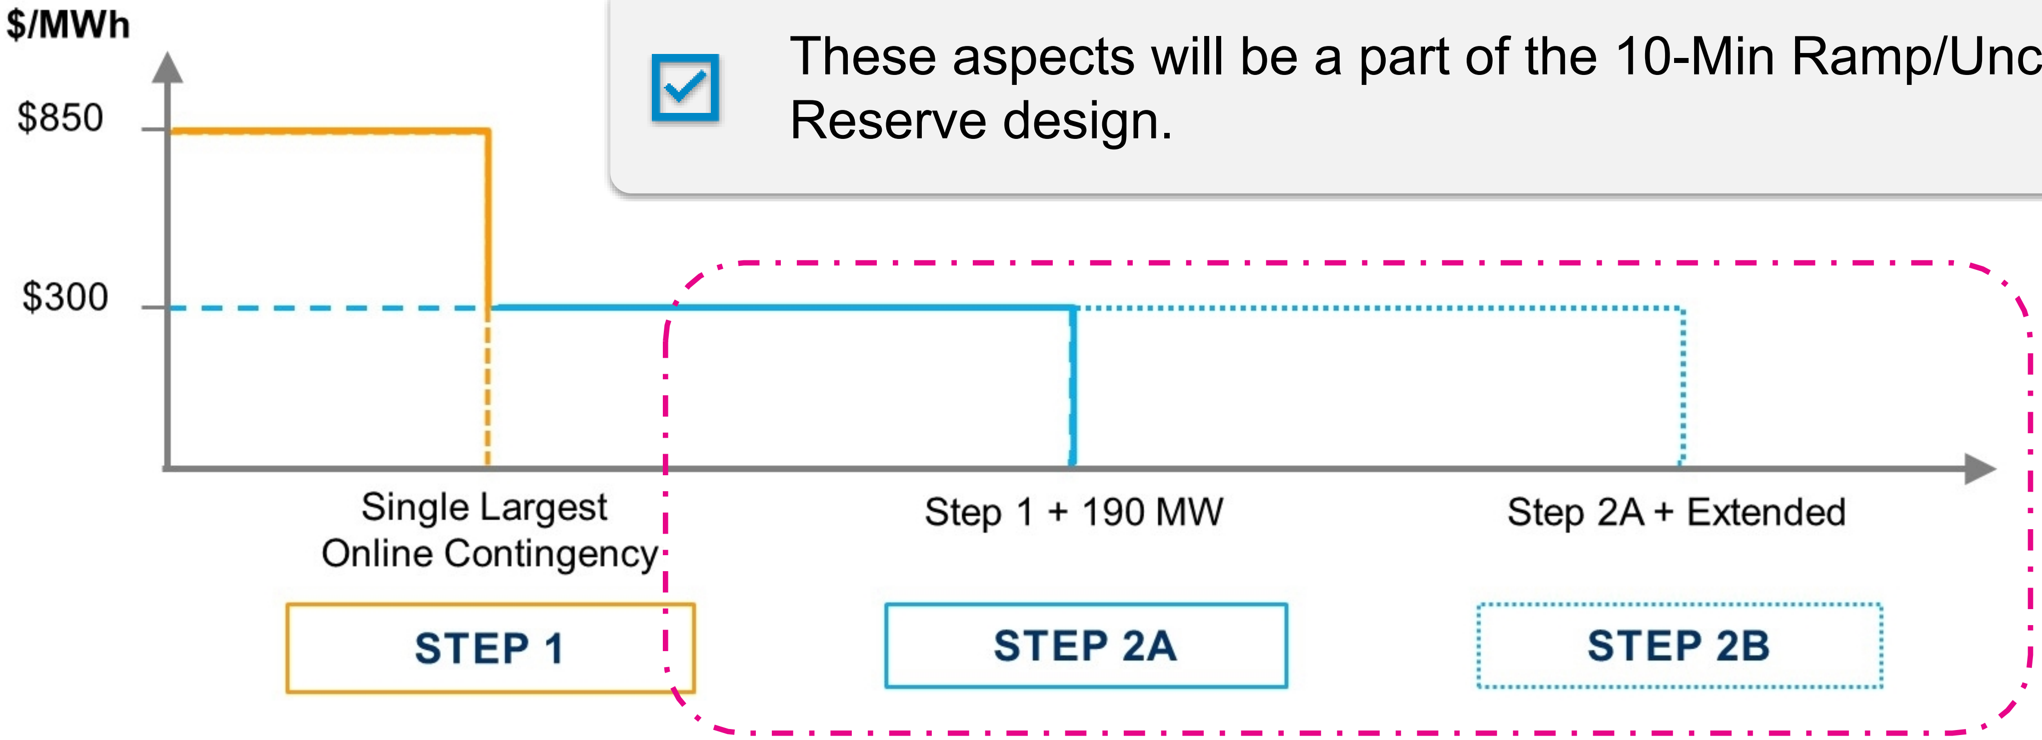

Other SR/PR Demand Curve Considerations

- SR and PR Demand Curves are the same today, but we may want to consider changing that in the future.

- The impact of product nesting on price formation means that we should consider how the curves represent the market's willingness to pay in a situation where we are short both products.

SR/PR Demand Curve Shape

PJM's current thinking is that unlike in our current ORDCs, some value differentiation in the penalty factor for reserves *below* the Reliability Requirement could more accurately reflect the value of reserves under different reliability scenarios.

- The Reliability Requirement represents PJM's NERC-mandated reliability obligation and ensures that PJM can recover from a unit loss. This would set the right-most point on the curve.
- Below the Reliability Requirement, the SR/PR demand shapes might rise until 50% of the requirement. This is consistent with the level of reserves we can request through the Northeast Power Coordinating Council reserve sharing agreement.
- Any reserves carried *above* the Reliability Requirement to manage net-load forecast and other uncertainties would be handled by other products, such as the 10-Min Ramp/Uncertainty Reserves.

SR/PR Demand Curves Conceptual Design

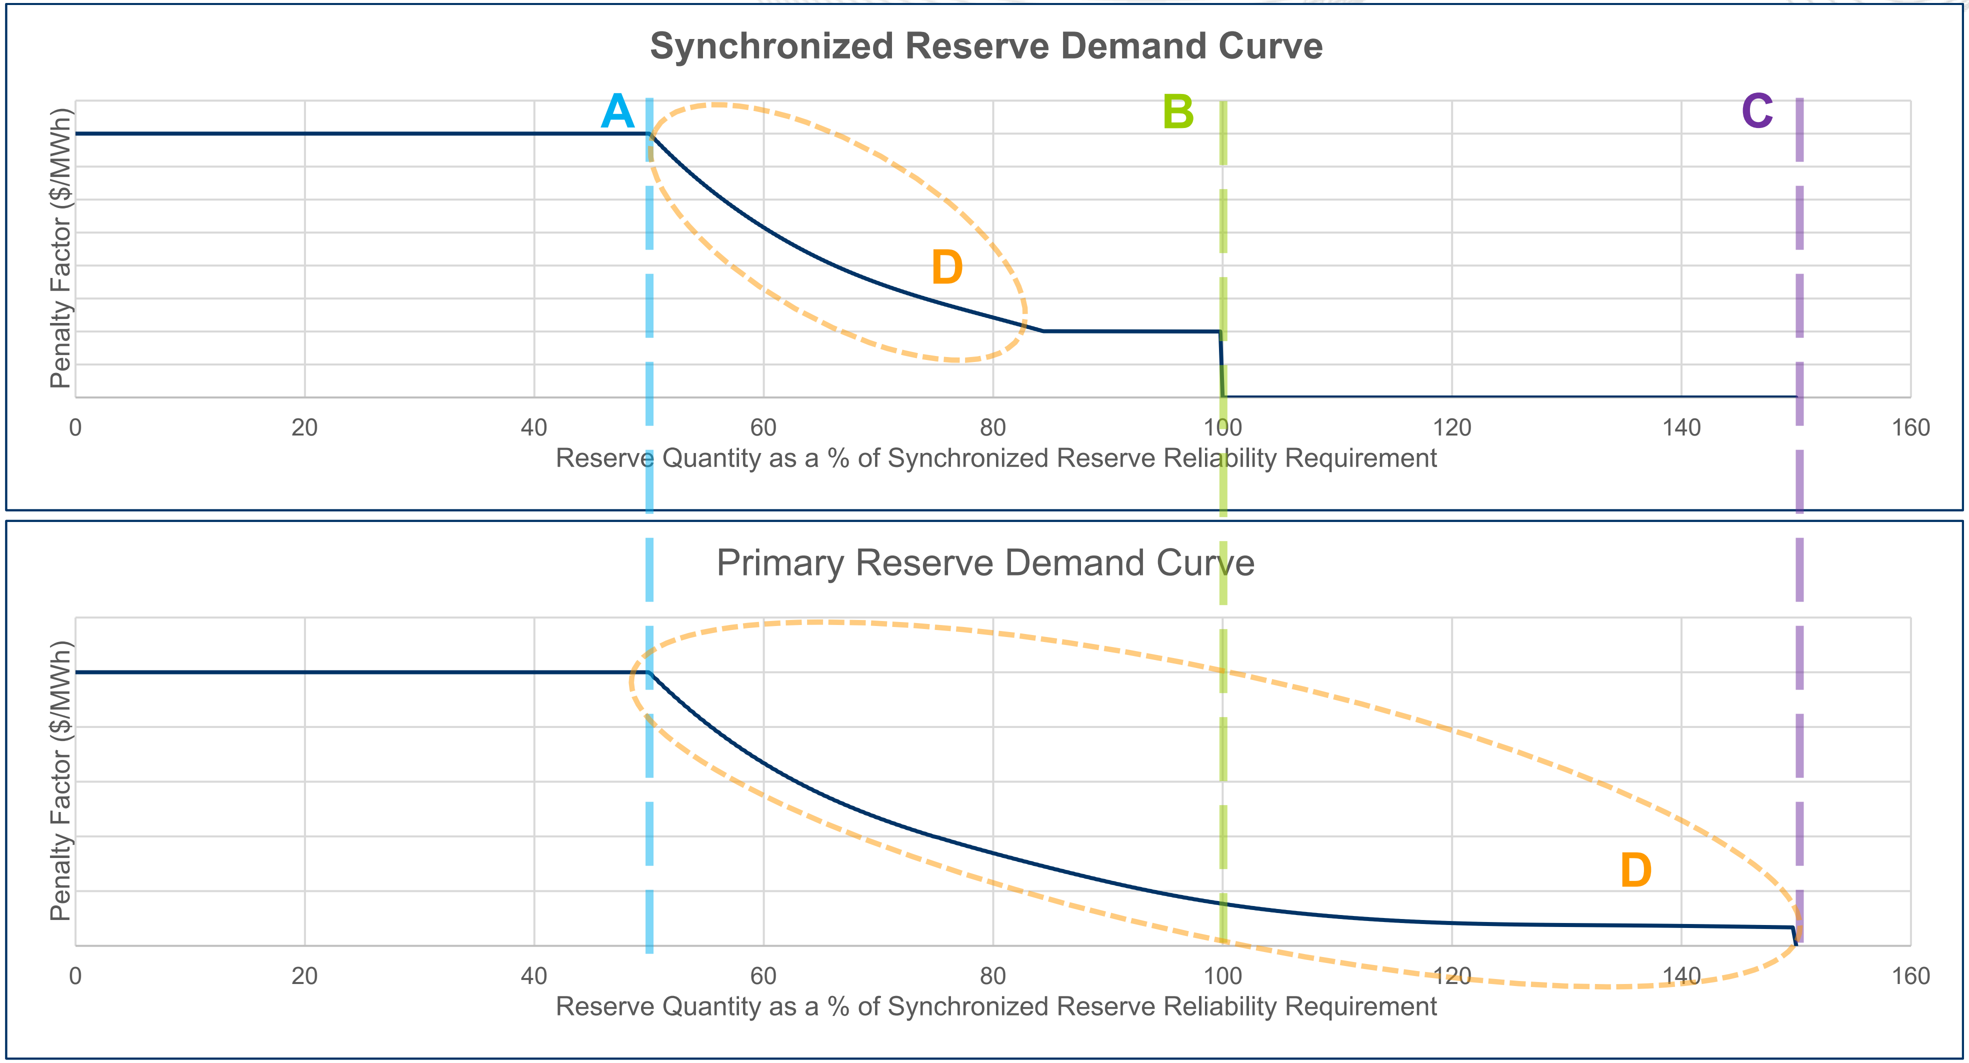

A 50% of the SR Requirement and 33% of the PR Requirement.

B 100% of the SR Requirement and 66% of the PR Requirement.

C 100% of the PR Requirement.

D Downward slope defined by the uncertainty distribution of historical generation forced outages.

PR/SR Demand Curve Shape: Uncertainty Calculation for the Downward Slope

Why consider using an uncertainty analysis to inform the demand curve shapes?

- As reserve quantities decrease, the probability of load shed increases. Historical operational data can provide insight into these probabilities, which in turn inform the value of reserves based on their value in mitigating reliability risk.
- For SR and PR, which is used to manage the risk of unit loss, the probability that we will need to shed firm load is based on the likelihood that the system experiences a generator forced outage without sufficient SR and PR to replace the MWs lost.
- Therefore, an uncertainty analysis looking at historical generator forced outage rates may be an appropriate way to define the shape of the SR/PR demand curves.

PR/SR Demand Curve Shape: Uncertainty Calculation for the Downward Slope

How might the uncertainty analysis be performed?

- Uncertainties would be calculated based on historical Generation Outage reporting.
- We would only consider outages starting within a limited time of the reporting (e.g., outages that are reported within an hour of the outage occurring.)
- An uncertainty distribution would be fit on that data using kernel density estimation, and the loss of load probability would be calculated as 1 – the cumulative distribution function (1 – CDF). This would define the downward slope of the curve (i.e., the shape). This curve would be anchored to a set penalty factor, for example, one set based on an analysis of resource commitment costs.

Other Market Design Considerations alongside SR/PR Demand Curve Reforms

Today, we have administrative price capping on our reserve prices if we hit shortage pricing for multiple products in the RTO and reserve sub-zone. As we consider changes to our locational procurement and ORDCs we will need to revisit this.

Our ORDCs should be coherent with our transmission limit penalty factors in the optimization software such that the optimization engines appropriately evaluate the trade-off between reserve shortage and managing network constraints.

10-Minute RUR MW Quantities

10-Min Ramp/Uncertainty Reserves (10-Min RUR) would be cleared to manage net-load forecast uncertainty and to ensure that the system is positioned to meet expected net-load ramping needs in upcoming intervals.

- Near-term load, wind and solar forecast error (e.g., 30 minutes ahead of the target time).
- Forecasted net-load ramp for intervals beyond the current target interval (e.g., the next 10 minutes beyond the target time).
- Both up and down 10-Min Ramp/Uncertainty Reserve products would be defined and the factors dictating the reserve quantities would be the same. However, the actual quantities could be different depending on system need.

10-Minute RUR Penalty Factors

10-Min RUR would be cleared in a co-optimization with energy and other reserve products by market clearing engines which consider the trade-off between going short reserves at the cost of the penalty factor or incurring the cost of clearing reserves.

- In part, the value of procuring 10-Min RUR would be to mitigate the risk of Synchronized Reserve shortages in future intervals.
- This may reduce price volatility while recognizing the value of this flexibility to real-time operations and promoting market efficiency.
- The greatest penalty factor for 10-Min RUR is therefore lower than the lowest penalty factor for Synchronized Reserves, and these curves should fit together coherently.

10-Minute RUR Demand Curve Shape

Procurement of 10-Minute Ramp/Uncertainty Reserves would be driven in part by uncertainties associated with net-load forecast error, and the shape of the 10-Min RUR Demand Curve could naturally be based on this uncertainty distribution.

- ☑ Probabilistic analysis could be based on 10-minute net-load ramps plus the 10- to 30-minute net-load forecast error.
- ☑ This analysis would be bifurcated into up and down ramping/uncertainty needs to inform up and down products.
- ☑ The uncertainty distribution fit on these data would generate a curve that describes the probability of going short reserves based on historical operations.

10-Minute RUR Demand Curve Shape Conceptual Design

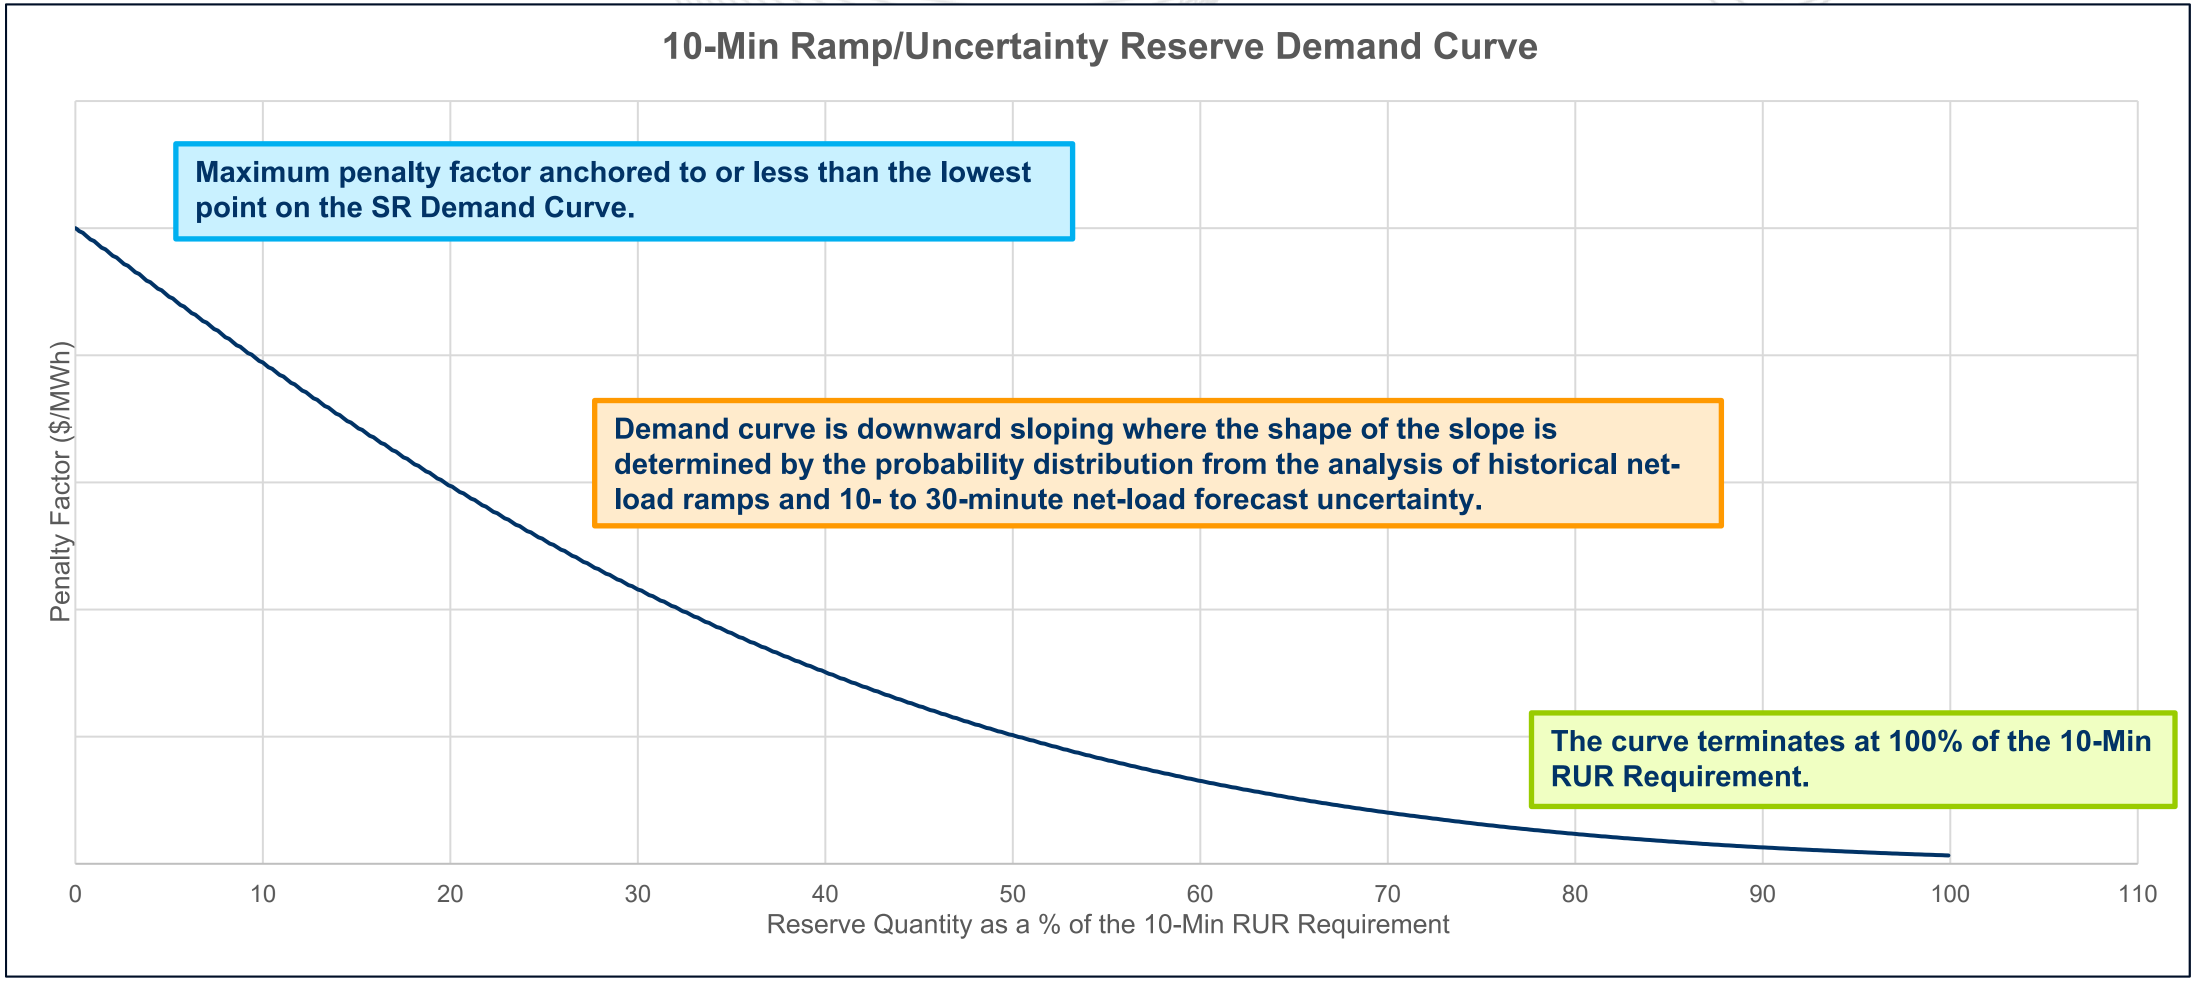

10-Min RUR Demand Curve Shape: Uncertainty Calculation for the Downward Slope

Why consider using an uncertainty analysis to inform the demand curve shapes?

- As 10-Min Ramp/Uncertainty reserve quantities decrease, the probability of going short Synchronized Reserves in future intervals increases.
- The probability of going short reserves is driven by net-load forecast uncertainty.
- Historical operational data can provide insight into these probabilities, which in turn inform the value of reserves based on their value in mitigating reliability risk.

10-Min RUR Demand Curve Shape: Uncertainty Calculation for the Downward Slope

How might the uncertainty analysis be performed?

- For 10-Min RUR, which would be used to manage net-load forecast uncertainty and expected net-load ramping needs in future intervals, the uncertainty analysis would be based on historical 10-minute net-load ramping events plus the 10- through 30-minute net-load forecast error.

- These data would be bifurcated to consider only positive values when looking at up ramping needs, and only negative values for down ramping needs.

- The data would be normalized by some quantile of the data set (e.g., 98-quantile), and once again the distribution fit would be used to set the curve as 1 – CDF.

- This curve would then be scaled by the penalty factor to anchor the curve at the left-most side of the curve.

30-Minute Reserve MW Quantities

30-Min Reserves (30-Min) would be cleared to manage net-load forecast uncertainty and to ensure that the system has sufficient Secondary Reserves to replace our Primary Reserves in the event of a unit loss.

- Load, wind and solar forecast error (e.g., 2-hours ahead of the target time).
- Replacement reserves for the largest contingency at any given time.

30-Minute Reserve Penalty Factors

30-Min Reserves would be cleared in a co-optimization with energy and other reserve products by market clearing engines which consider the trade-off between going short reserves at the cost of the penalty factor or incurring the cost of clearing reserves.

- As discussed with the SR and PR penalty factors, the 30-Minute Reserve penalty factors could be informed by resource commitment costs.
- The highest penalty factor for 30-Minute Reserves should be lower than the lowest penalty factor for Synchronized Reserves as the market optimizations should always prioritize maintaining our minimum levels of Synchronized Reserves.

30-Minute Reserve Demand Curve Shape

Procurement of 30-Min Reserves would be driven in part by the risk of unit loss (and therefore the need to replace deployed Primary Reserves) and net-load forecast error.

- An analysis of the uncertainty distribution of net-load forecast errors could inform the shape of the 30-Min Demand Curve.
- It will be important that the 30-Min Demand Curve fits coherently with not only the SR and PR Demand Curves but also with the 10-Min RUR Demand Curve to ensure that the optimization engines appropriately evaluate these trade-offs in reserve procurement.

The Role of Dynamic Reserve Requirements

- PJM’s position is that reserve requirements in the markets should align with operational requirements and accurately reflect system reliability needs.
- Given that reserve needs are predicated on system conditions, reserve quantities should accurately reflect those drivers at any given time.
- For example, on high load or high-risk days, PJM may need to carry more reserves, which would increase the MW requirements defined on the X-axis of the respective demand curve.
- As such, the ORDC shapes would be defined as a function of the % of the reserve requirement as shown in previous conceptual graphics.

Contact

Facilitator:
Lisa Morelli, Lisa.Morelli@pjm.com

Secretary:
Amanda Egan, Amanda.Egan@pjm.com

SME/Presenter:
Emily Barrett, Emily.Barrett@pjm.com

Operating Reserve Demand Curves

Member Hotline

(610) 666 – 8980

(866) 400 – 8980

custsvc@pjm.com

Acronyms

Acronym	Term & Definition
ORDC	**Operating Reserve Demand Curves** are administratively determined curves that establish a relationship between the quantity of reserves available and the markets maximum price or willingness to pay for a reserve service.
SR	**Synchronized Reserves** are reserves provided by resources that are synchronized to the grid and can respond within 10 minutes.
PR	**Primary Reserves** are reserves provided by resources that are either synchronized or not synchronized to the grid and can respond within 10 minutes.
RUR	**Ramping/Uncertainty Reserves** are reserves that would be procured to manage forecasted ramp and uncertainty operational flexibility needs.
MW	A **Megawatt** is a unit of power equaling one million watts (1 MW = 1,000,000 watts) or one thousand kilowatts (1 MW = 1,000 KW).


Protect the Power Grid
Think Before You Click!
Be alert to malicious phishing emails.
pjm
Report suspicious email activity to PJM.
(610) 666-2244 / it_ops_ctr_shift@pjm.com

Appendix

ORDCs Proposed Under Reserve Price Formation