Summary

The purpose of this study was to evaluate if more flexible dispatch of demand resources would relax the demand response saturation issues identified by PJM in its May 2010 *Demand Resource Saturation Analysis* report.

Background

A demand resource acting as an RPM capacity resource may be called upon up to 10 times per year, for a maximum duration of six hour per event. With the increasing amount of demand response clearing in recent RPM auctions, concerns have risen that either of those parameters could reduce the effectiveness of demand response as a capacity resource: either the demand response would be needed more than 10 times in a year, or that calling on demand response would merely shift the daily peak to hours outside those in which demand response is committed.

To evaluate those concerns, PJM planning staff performed an analysis which found that as demand response rose past 8.5% of peak load, it would be needed more than 10 times per year with unacceptable frequency. Additionally, the study found that with DR at 8.5% of peak load, the maximum duration of events would need to be increased to 10 hours to ensure that daily peaks are reduced, not merely shifted.

These findings were, in part, based on the assumption that whenever any demand response is needed, all available demand response is called:

> Each DR invocation is counted as one event, regardless of the amount by which reserves drop below the 1,300 MW margin. This assumption is consistent with PJM operating practice and the practical reality of emergency conditions. PJM Operations calls for DR several hours in advance of the actual need when the peak load for that day is not known. Typically, PJM Operations does not have the need to call for the amount of DR that would restore reserves exactly to the 1,300 MW margin. Rather, on a typical PJM emergency day, all the DR in the affected area is invoked.
>
> *(PJM Report, page 7)*

In this study, we present an analysis that argues that relaxing that assumption and allowing demand response to be dispatched in smaller increments will significantly increase the amount of demand response that can be used before the 10 days/6 hours constraints become reliability issues.

Ten Interruption Requirement

To evaluate the effects of more flexible dispatch on the ten interruption requirement, we first duplicated PJM's analysis as described on pages 3-5 of the report.[1] Our results, though not identical to PJM's, were close enough to demonstrate that we had successfully reproduced PJM's methods. In the graph below, our results for the base case are superimposed on the graph provided by PJM (report page 7):

[1] Thanks to Tom Falin of PJM for providing the data needed for this.

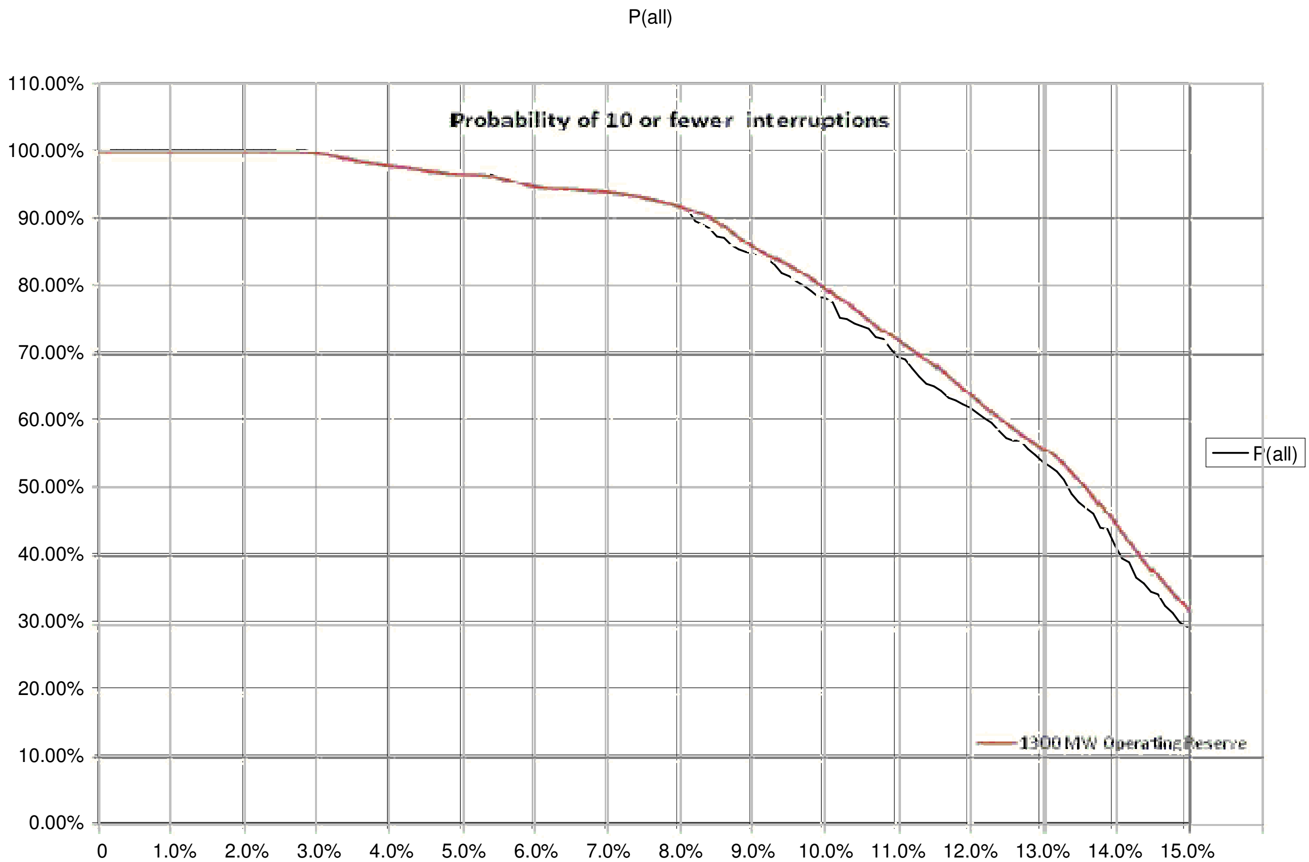

PJM's results are the smooth rust colored line; CPower's are the irregular black line. We have not identified the source of the difference, but believe it is small enough that our conclusions are generally valid.

To determine the effects of different dispatch strategies, we needed to modify PJM's approach. In their analysis, PJM determines the amount of required capacity as scenario peak load plus operating reserve requirement. This requirement is compared to a probability curve of the amount of available generation (the "PRISM Cumulative Probability Table") to determine the probability that demand response needs to be dispatched in any given day.

Because we are concerned not only with if demand response is needed, but how much is needed, we had to extend this method. Based on the same generation availably curve provided by PJM, we determined the expected amount of available generation as the probability weighted average available capacity from the PRISM Cumulative Probability Table. (In practice, this value turns out to be within 100MW of the simple 50th percentile value from the table).

The amount of demand response required on each scenario day, we took the scenario day's peak load (see PJM study, steps 1 and 2 on pages 3-4 for details of how those are determined), added in the 1300MW reserve margin, and subtracted the expected amount of available generation:

Expected DR needed = Peak + Operating Reserve – Expected Gen Available

This allowed us to create a chart, analogous to Figure 1 from PJM's report, showing the MW of DR required on each scenario day:

MW of Demand Response Needed					
Scenario	CP Day 1	CP Day 2	CP Day 3	...	CP Day 20
A1917	0	0	0		0
. . .					
A2006	14149.1	14149.1	11765		0
. . .					
. . .					
M2007	13483	7513	4170		0

For each day, the amount of DR needed was capped at the total amount of DR procured, 14,149.1 MW in the example above.

Next, we identified four dispatch strategies:

1. Dispatch in two blocks, each with 50% of total available DR
2. Dispatch DR in four blocks, each with 25% of total available.
3. Dispatch in 500MW blocks
4. "Perfect dispatch" in 1MW increments

For each scenario day, the amount of DR dispatch was determined as the amount needed, rounded up to the next largest block size, but still capped at the total amount of available DR. For example, if the block size was 500MW, any need from 1 to 500MW would result in 500MW of DR being called, a need for 501-1000MW would call for 1000MW, and so on. The total MW-days[2] of DR required was summed for all 20 peak days of each scenario, and a scenario was deemed saturated if the number of

[2] Using the term MW-day somewhat loosely to mean calling on 1MW of DR for one event; this is not the same as 24 MW-hours.

MW-days needed exceed 10 times the amount of DR procured. The change of DR saturation was then calculated as the percent of scenarios that were saturated. Our results are presented in the graph below.

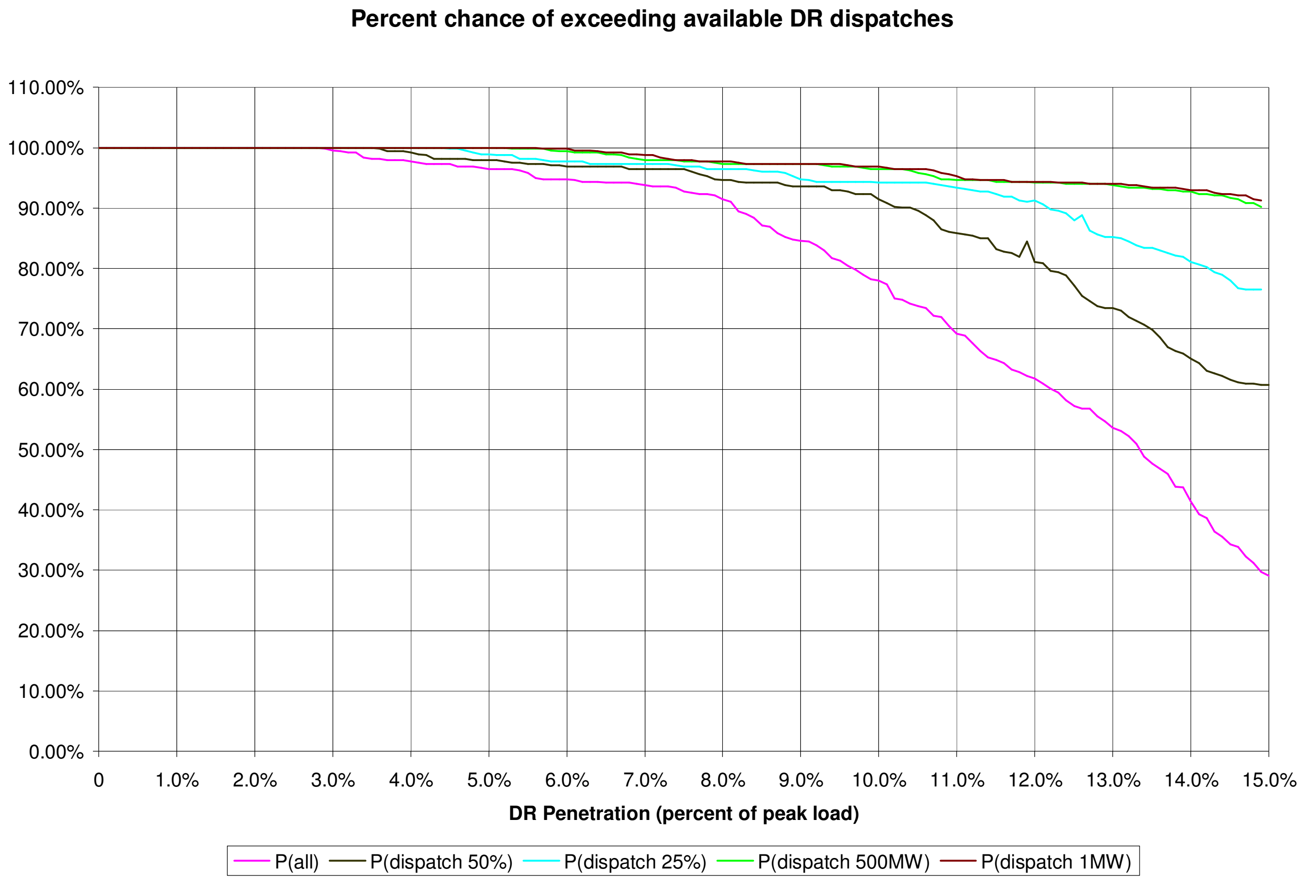

From this, we draw several conclusions:

1. Even dispatching DR in two batches each containing 50% of available demand response significantly increases the level at which saturation occurs. Using PJM's criteria that a 10% chance of DR saturation is acceptable (report page 8), we find that DR can reach 10.4% of peak load.[3]
2. Similarly, dispatching DR in four batches raises the allowable amount to 12.2% of peak, and dispatching in batches of 500MW raises the allowable amount to over 15%.
3. There is no need to strive for perfect dispatch; the saturation level when DR is perfectly dispatched is virtually the same as when it is dispatched in 500MW blocks.

[3] In two points on the above graph, the saturation level decreases as the amount of DR increases. This counterintuitive result is an artifact of the batching strategy, an only appears when the dispatch batch size is tied to the total amount procured. For example, since increasing the amount of DR also increases the size of a 50%, occasionally the amount of DR required on some days will change from slightly over 50% to slightly under. This means that the day's needs can be met with about half as many MW-days of DR, reducing the overall saturation level. Those spikes are arbitrary, very sensitive to deployment strategies, and should not be considered significant

Effects of Higher Installed Reserve Margins

Flexible dispatch of demand resources also increases the value gained from higher reserve margins. As the reserve margin grows, little additional demand response is needed on peak days (indeed, additional DR is only needed to meet peaks that would have caused reserve margins to fall below 1300MW under 15.3% IRM). Instead, the extra demand response is deployed on the secondary peak days, increasing the number of days DR can be called without exhausting the resource. Essentially, purchasing DR beyond the 15.3% IRM results in more days of deployment, spreading the demand response wider rather than deeper. The net effect of this is that the amount of allowable DR while maintaining a fixed chance of saturation increases by more than 1MW for each MW procured beyond the 15.3% IRM.

To demonstrate this analytically, we raised the reserve margin and found the highest level of demand response (as a percent of peak load) that would meet the "no more than 10% chance of saturation" criteria. For the 25% and 500MW dispatch scenarios, the results are:

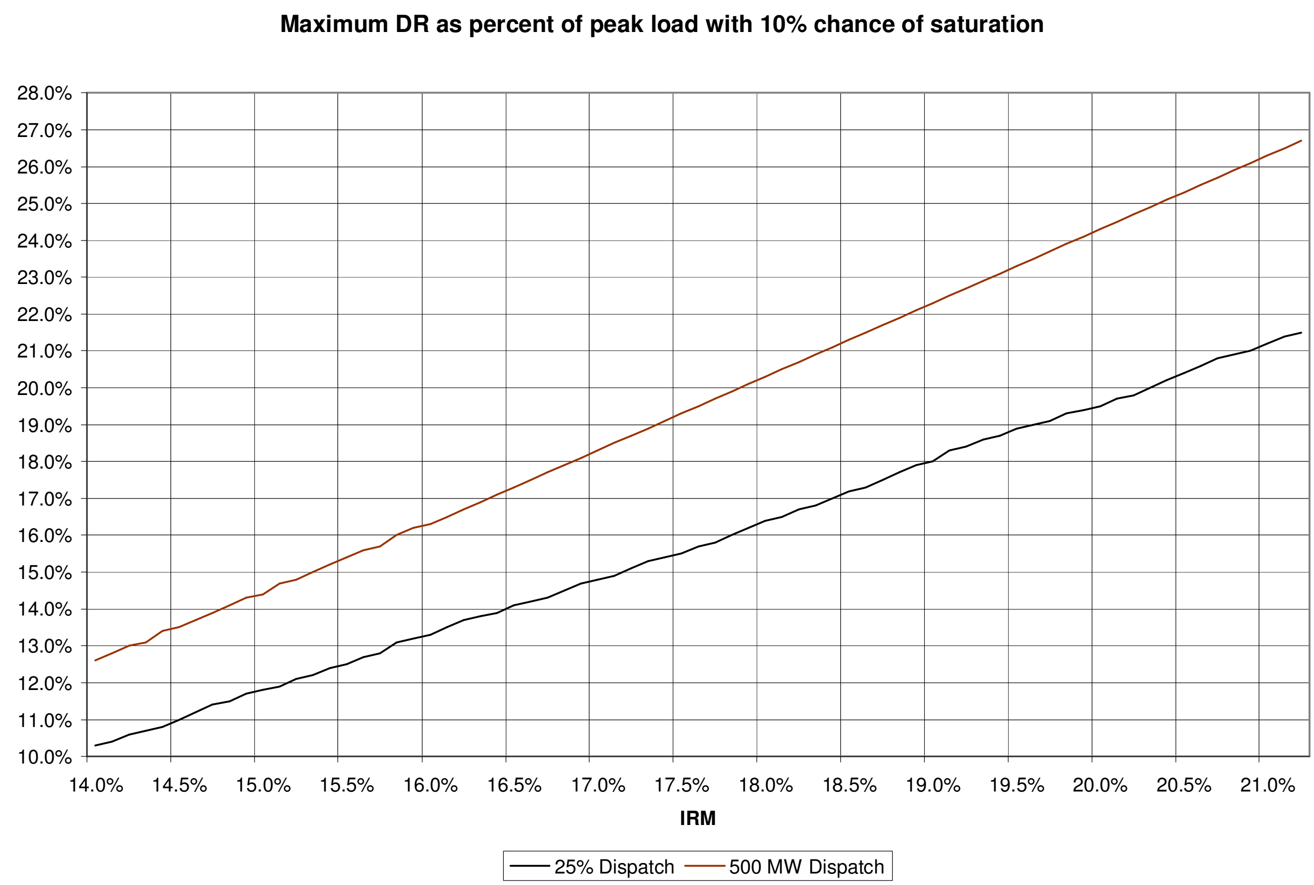

Given the reserve margins that RPM Base Residual Auctions have historically cleared at, and assuming PJM implements the 500MW dispatch option, the allowable amount of demand response rises to 20% or more of system peak load.

Six Hour Event Requirement
In the second part of their saturation analysis, PJM found that the maximum event duration would have to increase from 6 to 10 hours in order to avoid shifting the daily peak to non-event hours, assuming DR penetration levels of 8.5% of peak load.

We believe that it is important to distinguish between event duration at the system level and event duration for each individual DR resource. While we do not challenge the fact that events at the system level will start earlier and finish later as DR penetration increases, we believe that introducing more flexibility in how in DR is dispatched can help limit the event duration required of individual DR resources.

With flexible dispatch, DR saturation is reached when the amount of DR kilowatt-hours needed over the event duration at the system level equals the amount of all DR kilowatt-hours available, taking into account the maximum dispatch duration allowed for each DR resource. Essentially, this involves phasing in DR progressively so that the load curve flattens out at for the duration of the event.

In order to mirror PJM's analysis, we calculated DR saturation levels for each of the 10 load days used by PJM in their analysis. These 10 load days include 5 annual peak days and 5 days close to the annual peak during 2005-2009.

For each load day, we calculated DR KWh available using:
*DR KWh available = Max. Dispatch Duration * % DR Penetration * Unrestricted Peak*

To obtain DR kilowatt-hours needed, we added up the level of DR KW needed each hour in order for that hour's restricted load to remain at or below the daily peak with full DR dispatch. So for a given hour h, where Unrestricted Load > Daily Peak with DR:

DR KW needed = Unrestricted Hourly Peak– Daily Peak With DR
where *Daily Peak with DR = (1- % DR Penetration) * Unrestricted Peak*

Comparing the cumulative DR kilowatt-hours needed for each event day with the actual DR kilowatt-hours available during that day, we were able to solve for the level of DR saturation for a given maximum dispatch duration.

ILLUSTRATION OF HOURLY DISPATCH – 01/08/2006

- Hourly DR dispatch (blue columns) reduces peak over 9 hours (in green) even though no individual DR resource is called for more than 6 hours
- PJM's base case is in red: all DR resources are called for 10 hours
- In this example, DR saturation with 500MW dispatch increments is reached at 6.7% of peak load.

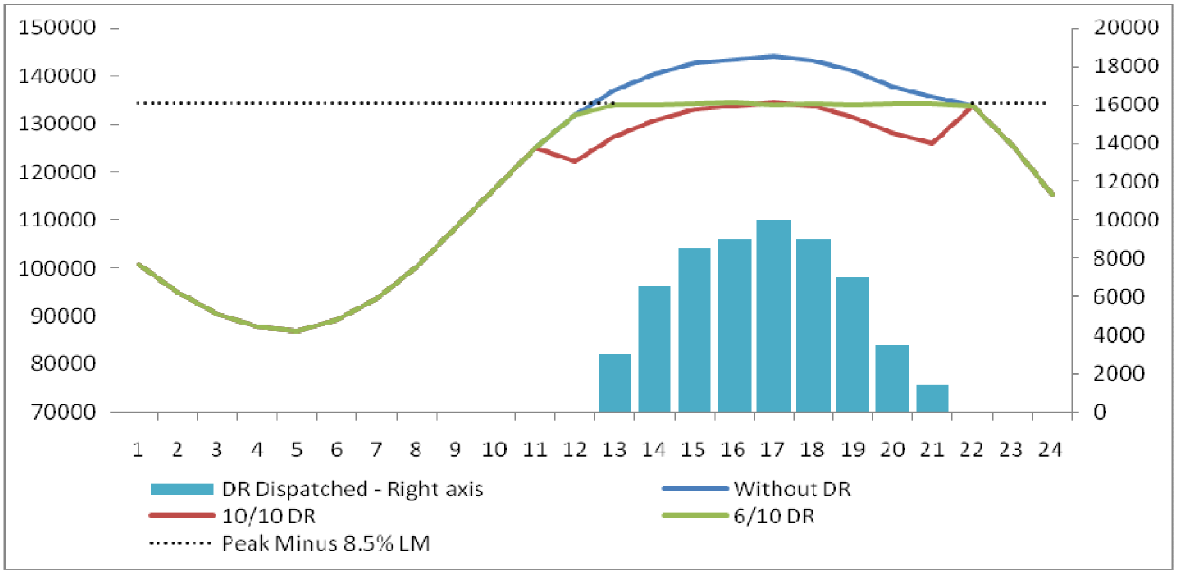

- If DR penetration is increased over the saturation level, the daily peak is shifted into the latter hours of the day, once all DR resources have reached their 6hr maximum dispatch duration.
- In this example, with DR penetration of 8.7% (2% over saturation) shifts the peak to after 8pm. This assumes that there is no voluntary curtailment beyond the 6 hour limit. The implication is that DR penetration needs to be several percentage points above saturation level for there to be a significant shift in peak load.

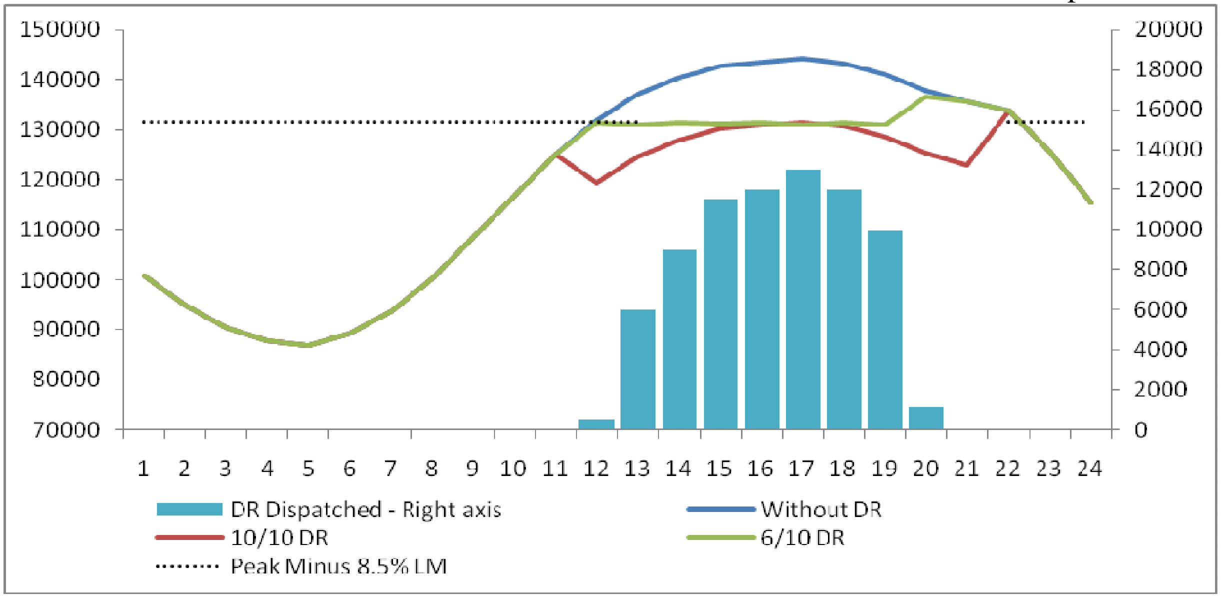

We evaluated six dispatch scenarios across the 10 load days selected by PJM:

1. Maximum dispatch duration of 6 hours; perfect dispatch in 1MW increments
2. 7 hours / 1MW
3. 10 hours / 1 MW
4. 6 hours; dispatch in 500MW increments
5. 7 hours / 500MW
6. 10 hours / 500MW

DR SATURATION LEVELS BY SCENARIO, WITH FLEXIBLE DISPATCH

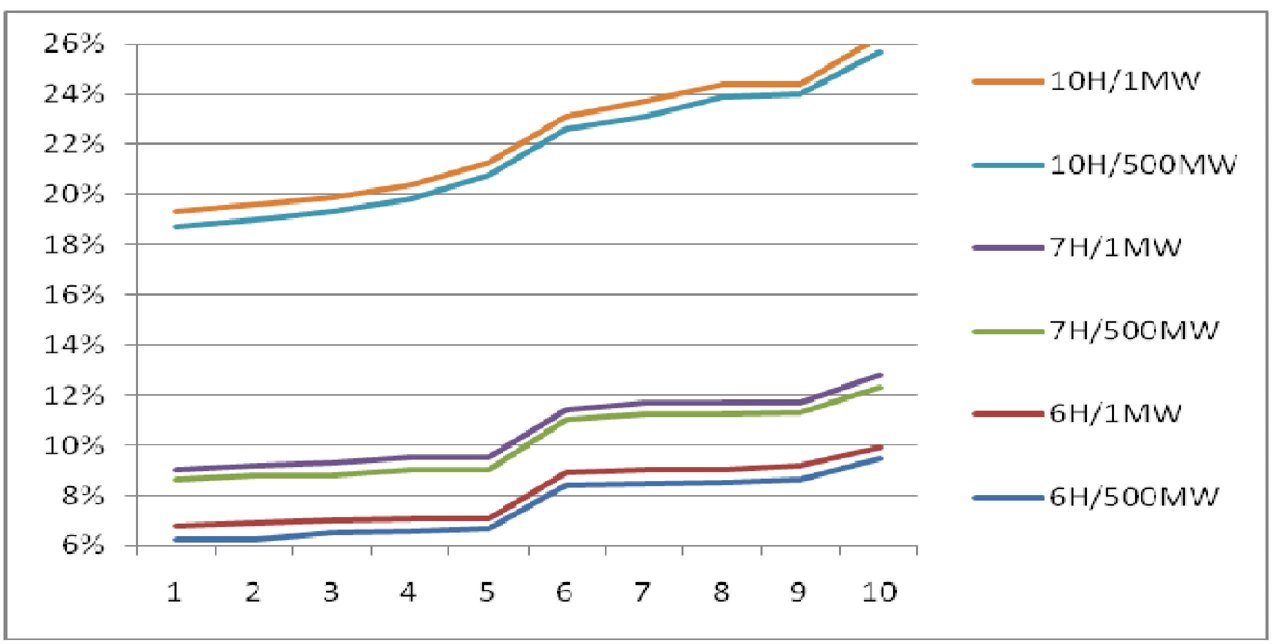

Summary of Results

With perfect dispatch in 1MW increments:

- With the current 6 hour maximum dispatch duration, DR reaches saturation at a penetration level of around 7% (range: 6.8%-9.9%).
- If dispatch duration is increased to 7 hours, DR's penetration can be increased to just over 9% before saturation occurs (range 9%-12.8%)
- If dispatch duration is increased to 10 hours, matching PJM's base case scenario, saturation does not become an issue until DR penetration levels reach over 19.5% (range 19.3%-26.3%).

If DR is dispatched in 500MW increments, all DR penetration levels above are reduced by about 0.5%.

Our conclusions are as follows:

- Flexible DR dispatch significantly reduces the maximum dispatch duration needed compared to PJM's analysis.
- With flexible DR dispatch, a small one hour increase in maximum event duration, to 7 hours, could push back DR saturation concerns until DR reaches 8.5-9% penetration.
- Any extension in DR hours need not be for the full RPM commitment; lesser performance in hours beyond the sixth can eliminate the peak shifting. A performance requirement such as "100% of RPM commitment for six hours, 50% in the seventh hour and 25% in the eighth hour" may have equal reliability benefits and be more achievable than simply extending the number of hours required.
- Even if DR penetration exceeds saturation levels (which is likely to happen sooner if the maximum dispatch duration remains at 6 hours), the shift in the daily peak need not be significant unless saturation levels are exceeded by several percentage points.
- It is also unclear if load will return to its pre-curtailment levels following an event. There are some arguments that cooling loads experience a "bounceback" effect and return to higher than normal values after an event; however, CPower has heard from some of our industrial customers that long curtailments simply involve sending workers home early, in which case load would not return until the following day.

Conclusion

We caution that none of the people involved in this analysis have any formal engineering background or system planning experience, and we do not have access to the depth of data that PJM does. Our analysis also made no attempt to model distinct LDAs, let alone any other subtleties of the RPM process. This report is thus not intended as an authoritative analysis of the potential for reliable use of DR in RPM, but as a stimulus to further research.

However, we are confident in our core result, that dispatching DR in smaller blocks will significantly raise the level at which saturation occurs. Based on this, we request that, before implementing changes to the demand response product definition or setting caps on the amount of demand response that can clear in an RPM auction, PJM:

1. Perform its own analysis on the effects of more flexible dispatch on the saturation issue
2. Examine the effects of higher reserve margins on saturation
3. Ensure that any constraints placed on RPM clearing incorporate those results.