To: Dave Lawrence

From: Bernie Neenan

Date: March 8, 2001

Topic: More on CBLs for the DAM PRL Program

Below are some thoughts to follow-up on the discussions about CBLs for the PRL DAM program that were held on Wednesday, February 27th.

There appears to be strong support for some form of dynamic CBL because many believe that it would be a) easier to administer and b) more representative and therefore inherently fairer. In some customer situations both conditions attain. However, as the discussion below describes, under other circumstances one only applies, and there are situations where neither attains.

The dynamic CBL is easier to administer. The relative ease of administration depends upon how the CBL is to be used. A predetermined CBL is developed in advance and must meet both the customer's and the supplier's approval. Assembling, validating, and reconciling 8760 load profiles from historic hourly interval data to create a predetermined CBL is time consuming. Creating a predetermined CBL becomes even more tedious when there is no interval data available for a customer. In that case, the first task is to identify what profile to assign to the customer and then reconcile the profile's individual hourly values to the available metered monthly data, the demand and energy values. Moreover, in some cases the customer's circumstances have changed requiring some additional adjustments to the CBL before it is mutually acceptable to the customer and program underwriter.

Finally, the CBL must be date-mapped; the historic days must be projected on the calendar year. Utilities with contemporary load research tools and systems report that preparing a CBL, not counting time spent reviewing it with the customer, requires 1-4 hours, depending on the quality of the interval data and the number and extent of adjustments resulting from the customer review process. The preparation time for CBLs in RTP is highly variable from customer to customer, due to the nature of the data available and because the predetermined CBL is developed in conjunction with the program marketing process, which in many cases involves considerable customer interaction and participation in the process.

Creating a dynamic CBL (DCBL) as a ten-day rolling hourly average values at the end of each month, using historical interval data that are already reconciled and validated, is simple and straight forward by comparison. If all CBL calculations are performed after-the-fact, then the dynamic approach offers enticing ease and simplicity.

However, when the DCBL is developed on the fly, every day, the relative advantage of the dynamic method is eroded. A customer who wants to have its CBL available when it prepares its daily PRL DAM bid must have access to its metered data for the prior day (readings up to midnight) before the 5:00 a.m. bid submission deadline. Provided that the metering authority serving the customer agrees, the technology is available for interrogating the meter's register, streaming the previous day's hourly values into a database, and appending it to previously collected values, from which the dynamic CBL can be calculated. Several vendors offer such technology, which ranges in price from few a hundred dollars to several thousands for sophisticated, real-time systems. Even with this technology in place, the customer undertakes some risks since the data it retrieves and utilizes for its bids are not the data that will ultimately be used to reconcile its performance, should the bid be accepted. Final resolution of the CBL comes only after the meter is read in due course and the results are reconciled following the procedures employed by the metering authority.

So, which is easier? The predetermined CBL requires an up-front investment of time by both the customer and the program underwriter. Moreover, underwriters who anticipate contacting a large number of customers would likely have to acquire or develop specific analysis tools to accommodate the CBL preparation needs. However, once the predetermined CBL has been prepared and accepted, there is no further work involved for either the customer or the underwriter. Customers have access to the exact hourly CBL values that will be used to measure program performance a day, week, a month and a year in advance, a feature that RTP participants value highly.

The dynamic CBL requires no analyses prior to the customer’s joining the program, which speeds up the marketing process. And, as long as the customer does not require accurate CBL values on an ongoing basis, preparing a dynamic CBL once a month from approved billing data is analytically a straightforward and relatively simple task for the program underwriter. However, if the customer demands accurate and timely CBL information so that it can prepare its daily curtailment bid, then maintaining simplicity from the customer’s perspective requires an investment by either the customer or the program underwriter. Some customers may already have in place equipment for interrogating the meter to acquire readings. Most however would have to have such equipment installed and maintained, which in itself creates a deterrent to subscription, either because of the cost or inconvenience, or both.

The Dynamic CBL is fairer. There is nothing inherently fair or correct about a dynamic CBL relative to a predetermined CBL. The dynamic CBL utilizes historic usage values to establish each day’s typical usage, just like the predetermined CBL. The difference is that the predetermined CBL value for each hour is the corresponding hourly usage level from the same day a year prior, while the dynamic CBL uses the average of the corresponding hour’s usage from ten days prior. Which one is fairer depends upon whether the customer’s ongoing daily usage profile is best typified by its most recent usage level or its past year’s usage level.

Lagged data structures like the dynamic CBL are commonly employed in forecasting situations where the variable of interest follows a systematic cycle of an upward trend followed by a downward trend, but its path along the way is subject to short-term fluctuations. When the trend is moving upward, the rolling average pushes the forecast upward at a rate that reflects an ongoing tread. Short-term fluctuations that result in values counter to the trend become dampened by the use of several past observations, most of which follow the current trend. The length of the lagged structure to use depends on the relative frequency and size of the occasional fluctuations that divert values from the trend. If the fluctuations are relatively large in size, then a longer lagged structure will reduce the distortion and better reflect the overall trend.
When the variable to be forecast is following a flat trend, i.e., there is no upward or downward direction, but it is subject to occasional up or down fluctuations in the value,

the lagged forecasting process is equally applicable. The lagged values, most of which are constant, temper the occasional variation from the trend and diminish its affect on the forecasting process.

A dynamic CBL is most representative for a customer whose hourly pattern of usage follows a very stable daily pattern but is subject to random fluctuations. As the hourly loads rise gradually from day to day, the lagged structure of the DCBL increases the next day's projected value for each hour according to the established upward trend. When the trend reverses, the downward pattern is captured, and the counter-trend fluctuations are dampened. Like any lag structure, the D CBL process produces the worst forecasts of the customer's likely usage at turning points, when the trend changes from upward to downward or vice-verse: the lagged values include observations from the previous trend and pull the forecast in the wrong direction. The longer the lagged structure, and the shorter the cycle time in the trend, the greater the turning point distortions.

A predetermined CBL can never match historic weather exactly to what is encountered during the year. The greater the customer's usage sensitive to weather to other random factors whose occurrence varies from year to year varies, the greater the disparity on those days between the customer's usage intentions, arising from contemporaneous conditions, and the CBL assigned to the hours of that day.

A dynamic lag however can adjust the CBL to reflect weather patterns as long as the lag structure mimics the weather pattern. Take the case of a customer whose load is stable except for weather effects, which cause variations. For example, if a one-day lag is employed in the CBL calculation, when a three-day heat wave hits the hypothetical customer with the highly weather sensitive load, say on a Monday. The Wednesday CBL, created Tuesday morning using Mondays metered usage (the one-day lag), would be adjusted upward to reflect Monday's incremental, weather induced, load and correctly portray what the customer might be expected to use on Wednesday. Curtailment bids should be measured from this elevated level of usage, not the typical, weather normal, level.

But, what about the DCBL for Tuesday and Thursday created from this on-day lag structure? Early Monday morning, the Tuesday CBL is derived from the previous Friday's usage, which was normal and therefore at a lower level. So, on Tuesday, the customer's DCBL is low relative to what it expects to use given the heat wave. That means that in order to perform to a curtailment obligation, it first must provide "free curtailment" from its actual operating level down to the biased (downward) DCBL level before it gets paid, meaning that it must actually drop more load than it gets paid for to avoid noncompliance penalties.

As described above, the Wednesday DCBL accurately portrays the situation, and the DCBL is fair. By Thursday the heat wave will have passed and the customer usage requirements will have decreased back to normal. But its DCBL, developed Wednesday morning, remains elevated because it is constructed from Tuesday's elevated level. For Thursday, the customer enjoys a DCBL biased upward. It can bid load that it will not need and therefore it can perform to its bid with less inconvenience and cost. The same holds true for Friday, constructed from Wednesday's high heat-induced load.

To recap, a one day lag in the hypothetical case constructed above for a highly weather sensitive customer results in one day where the CBL is biased downward, one day where it is spot on, and two days where it is biased upward. The customer has little incentive to bid the first day because of the windfall loss. On the second day it has the proper incentive. On the next two days it can over-bid with no inconvenience or outage cost because its CBL is biased upward. Even a small CBL bias can heavily influence the bidding strategy, especially when customers intend to bid only a small portion of their load for curtailment; the windfall gain or loss is proportionally larger.

Extending the lag to say, three days, and using a simple average of the values, provides some mitigation of the bias, because in effect no day is ever fully representative. When weather effects are of longer duration, and especially in climates where for weeks on end weather patterns persist, the longer lag structure performs better. It levels out the relatively minor peaks and valleys so that they should have only a small impact on the customer's bidding strategy. If the customer bids often and consistently, then the bias will balance out.

A ten-day average is being used in the PRL emergency program and its extension to the day-ahead program has been proposed by some as appropriate. For customer's whose load variations are weather related, this might provide a suitable means for weather normalization, especially, with respect to seasonal patterns (the ten-day average would distort short-term weather aberrations more than the three-day average).

However, this long lagged structure has a perverse impact on customers who have a very stable usage pattern that follows a weekly cycle. The figure below illustrates this

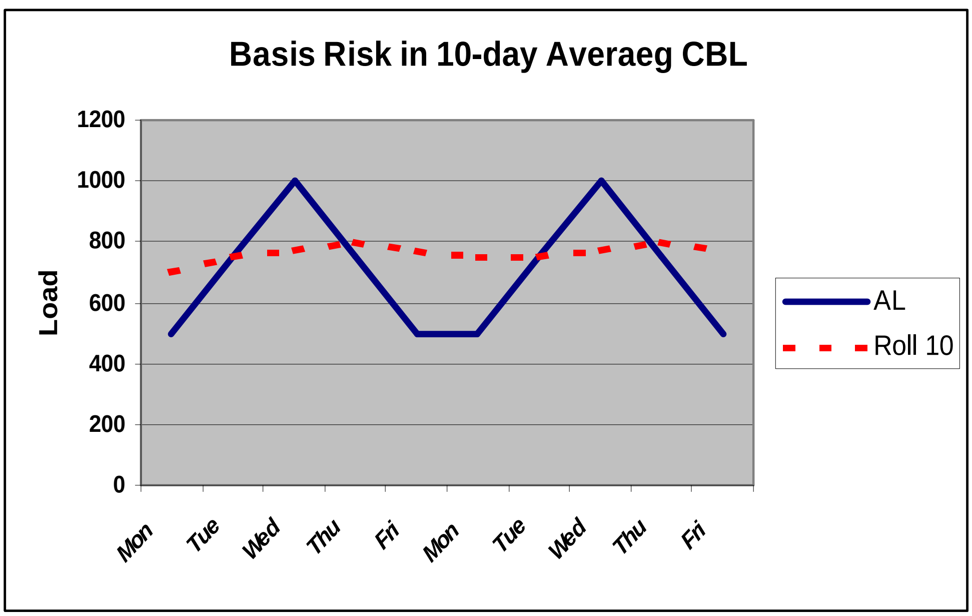

basis risk.
The solid curve represents the customer's weather independent, weekly pattern for one hour (e.g., 4:00 p.m.) that is repeated week-in and week -out. The hourly load portrayed builds from Monday through Wednesday and then descends back in the same pattern

until Friday's matches that of Monday. The dashed line shows the ten-day average DCBL. In only two days out of ten is the CBL representative. The rest of the days it is either high or low, exposing the customer to basis risk—either a windfall profit or loss if it bids on those days.

Averaging by day types might rectify this basis risk. For example, for an hour on any Monday, the value would be the average of that hour on the previous10 Mondays, not the most recent 10 weekdays. As long as the weekly pattern is not subject to seasonal or other periodic shifts, then this method would remove the gap between the CBL and the actual intended load. However, this extended averaging increases the cost and adds complexity. Is this still simple?

Another suggestion for better DCBL representation is to construct the DCBL from a recent historic period using not all of the days in the lag structure, but only selected ones. Two methods come to mind: use the highest *d* days or the lowest *d* days. For example, using the 10-day lag structure, the DCBL could be the average of the highest hourly values for the 10 days in the lagged sequence. This method exacerbates the bias and introduces opportunities for gaming.

The figure below illustrates the customer with the reoccurring load pattern but with its DCBL based on the average of the three highest hourly values from those recorded in the previous 10 weekdays. By choosing the most favorable days from its perspective, the customer has a CBL that equals or exceeds it actual intended usage eight days out of ten. Therefore it has incentive to bid excessively on a regular basis, and enjoy windfall profits from its bids.

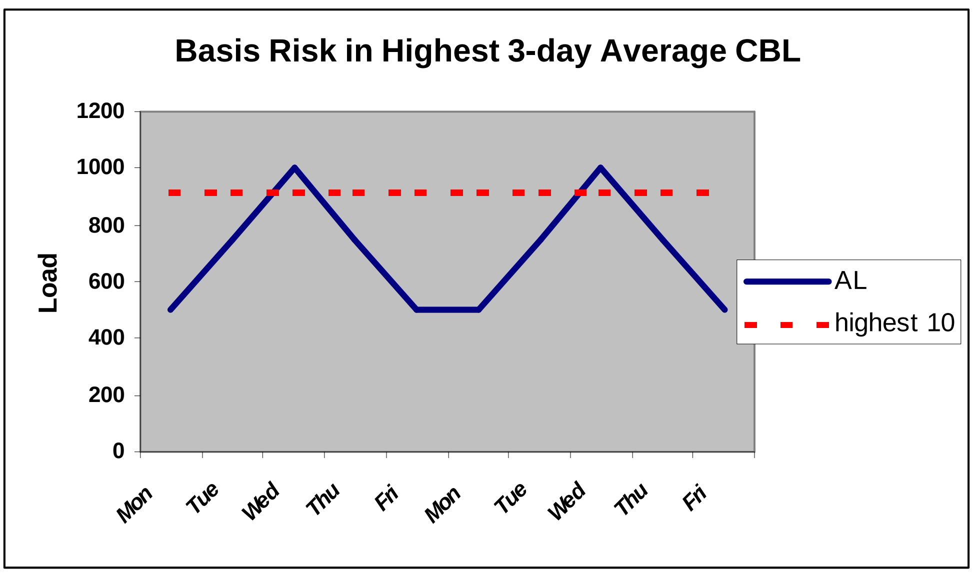

Moreover, a plausible gaming strategy is available under this CBL methodology. A customer could operate at full level for three days, and then shut down for three days and bid curtailment load at low prices on those days with no fear of incurring noncompliance penalties. While this case seems out of the ordinary for most customers, it demonstrates a strategy that customers could use to bid loads for curtailment whose impact on the system are negligible or nonexistent.

A Simplified Way to Develop the Predetermined CBL

Other methods have been utilized in similar situations to develop a predetermined CBL using a simplified methodology. The basic approach used was to limit curtailment bids to a specific daily window (e.g. from 8:00 a.m. to 8:00 p.m. weekdays), and then develop the CBL from the customer's metered maximum demand and energy consumption data used for billing. The CBL for each hour of the window was equal to the monthly maximum demand from the historical year (e.g., the previous year) multiplied by the customer's load factor for that month. The result is a flat CBL for each hour of the window.

In this methodology, the load factor adjustment de-rates the customer's maximum usage relative to its average usage. Customers whose load is the same every hour of every day would have a load factor of 1.0 and rightfully its CBL would be equal to its maximum demand. Customers whose loads rise and fall throughout the window would have a CBL that is less than their maximum demand. The figure below illustrates this methodology. The maximum demand is the highest value in the period 8:00 a.m. to 8:00 p.m., in this case 1250 kW.

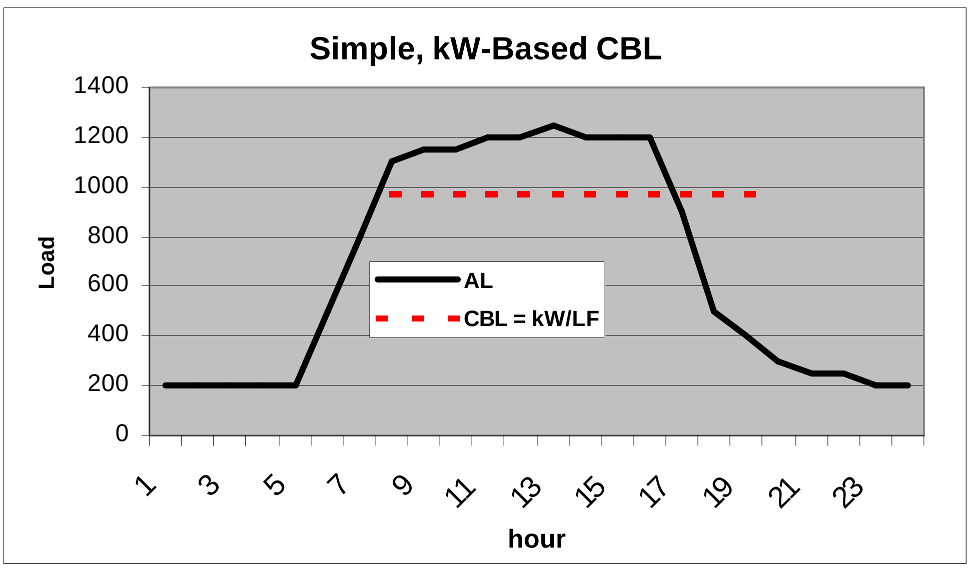

The load factor is defined as the maximum of the sum of the hourly energy usage levels during that time period divided by the product of the maximum demand and the number of hours in the window. In the example portrayed, the load factor is 78% over the 12-hour window. As the figure above illustrates, the customer's CBL is biased downward during the hours of the day, noon to 6:00 p.m., when the LBMP is likely to be highest and the customer's has the most load to bid. The CBL would be even more biased if the maximum demand used for billing were the non-coincident (with the CBL window) demand. The addition of the other twelve hours of the day plus weekend hours, whose loads are relatively low compared to the maximum demand, would reduce the CBL to 516 kW, and all but preclude the customer's participation, even if he has ample load management capability to participate and contribute during the mid-day hours.

A Max kW-based CBL provides a better representation if its usage is limited to a few hours of the day. Extending the example from above, if the bid window were restricted to the hours 10:00 a.m. to 5:00 p.m., the CBL based on the Max kW method for those seven hours would be over 1150 kW, which results in very little under representation during the narrower time window. The example in the figure below portrays this result. Under this CBL methodology, the customer must agree to limit its bids to the narrow window used to establish its CBL.

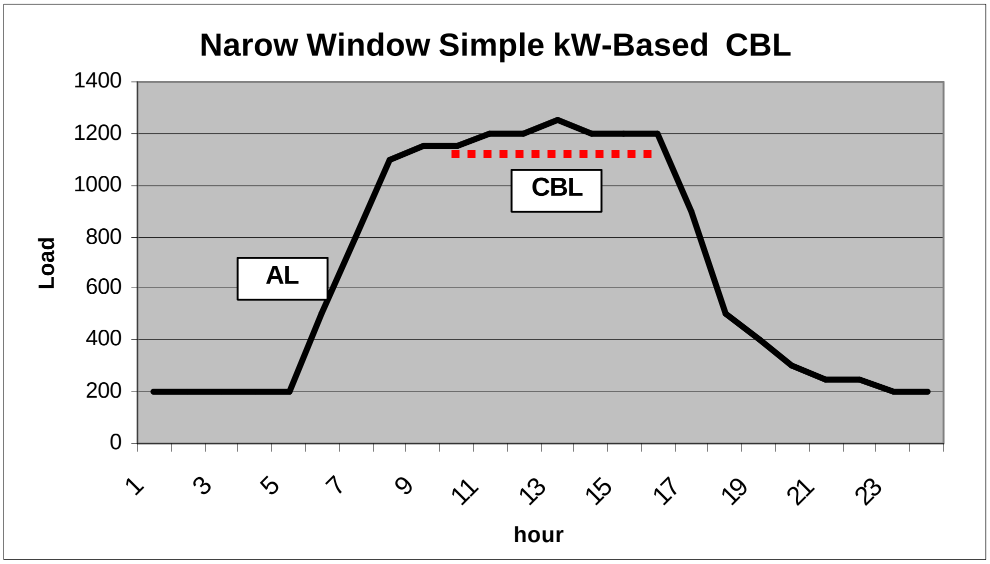

This alternative CBL method was employed in a load management program where the curtailment window was limited to the primary business hours of the day. Over 40 customers, many with maximum demands under 500 kW, participated using the simple kW-based CBL. Since customers' demands were measured as the non-coincident maximum kW, to avoid imposing under-representation *adjustment factors* were developed from class load data and applied to all customers in that category. The adjustment factors used were higher than the customer's load factors in many instances, better reflecting the customer's actual load curtailment capability in the window of interest and thereby enhancing program participation. Many markets are adopting methods like this to devise deemed load shapes that are assigned top customers for which there is no interval-metered data in order to settle hourly wholesale market transactions.

Other Mechanisms to Make the CBL Representative

Whether a predetermined or dynamic CBL methodology is employed, there will still be bias in how most customers' CBLs are established. Downward bias, the situation where the CBL under-represents the customer's intended usage, acts as a deterrent to program participation or limits the days a participant will offer a curtailment bid. Because customers have to over-curtail to meet their bid obligation, the actual value of the benefit is some portion of the LBMP. Customers can factor this over-compliance bias into their bids, which reduces their chances of having their bid accepted, or accept the fact that sometimes their returns from bidding into the program will be lower than others. Another remedy would be for the customer to revise its overall operating schedule in a way that reduces the CBL bias, although the cost of doing so might not be warranted in light of the expected benefits that would be realized.

Some have expressed concerns that in the case where the CBL is systematically and persistently upwardly biased, customers would bid this apparent curtailment capability everyday at relatively low prices thereby ensuring that their bids are accepted. Doing so would lock them into a benefit since their upwardly biased CBL will be above their actual and intended usage, resulting in an ongoing benefit stream for which there are no offsetting costs. One way to mitigate windfalls arising from systematic and persistent upward bias that overstates the customer's CBL is to place a floor on the level of curtailment bids. For example, the program could institute a minimum bid price of $100/MW. This will not prevent customers from bidding and realizing windfalls when prices are high. But, it would prevent them from realizing a gain from them under normal, low priced periods.

Another mitigating mechanism would be to limit curtailment bids to a predetermined threshold level, some percentage of the customer's typical usage. This might be accomplished by using the customer's maximum demand by month to establish a CBL basis, and then set a firm power level (FPL) that is some proportion of that basis, for example 50%. Each day the dynamic CBL would define the hourly CBL values and the customer could then bid up to the amount defined by the difference between the dynamic ceiling and the FPL. This limits the upward bias in benefits since the customer cannot bid any amount greater than the difference between the dynamic CBL value and FPL. If it shut down completely for a few days to take advantage of the upward bias in the dynamic CBL methodology (which would reflect the full operation level, or a large percentage thereof, for a few days), it would be limited to a bid of no more than 50% of the dynamic CBL.

Summary

Dynamic CBLs were frequently used in conventional interruptible tariffs to establish the level of benefit to be paid to participants and define their load curtailment obligation. The customer's metered maximum demand, or some function of the monthly maximum metered demands, was used to establish the participant's capacity basis. Payments were then made on the difference between that kW level and the participant's firm power level (FPL), which each participant selected based on its ability to meet the curtailment obligation. Simplicity was one reason this methodology was favored. Another was that this construction befitted a program designed to pay customers for the effective capacity they provided to the system. But, it was not without its fallacies.

Rate designers recognized that there was inherent bias in this formulation, since unless a demand for an interruption coincided with a time when the participant's usage was exactly equal to the maximum demand used for the payment basis, the participant would enjoy a degree of windfall benefit – it was paid for load it did not have to forego. To compensate, the benefit (a $/kW-year amount) was discounted by a factor that took into account the basis overrepresentation. Usually, the smaller the customer, the lower the discount factor, thereby reducing the benefit according to participant's load factor. This design was attractive to very high load factor customers, but less enticing to lower load factor customers who were paid lower benefits.

Predetermined CBLs have been used almost exclusively in the retail RTP programs implemented in the past 15 years. The RTP model focuses attention on energy transactions, unlike the interruptible model's capacity emphasis. In RTP, an hourly energy basis was required to define the customer's rights to firm power (under the tariff) and to value changes in usage from the CBL that would be subject to market price. A predetermined CBL was employed because it would provide the best representation, over the entire vector of 8,760 hours of the year, of what the customer would have used under the tariff, in the absence of the RTP program. Day in and day out it provided a relatively straightforward and equitable means for allowing customers to participate simultaneously in two markets with inherently different pricing rules.

Customers participating in RTP came to be appreciative of the predetermined CBL. First, while it required some effort to get the initial CBL developed, that effort was not wasted. It afforded the customer an opportunity to carefully evaluate its ability to manage load to avoid unintended consequences (paying high prices) and take advantage of the program's benefits (sell at high prices, buy at low prices). The CBL development process served as a screening device that made believers out of many skeptical customers and caused others to reject participation when they realized how little control they could exercise over their usage. Perhaps this is one reason for the low dropout rates reported for RTP programs.

Second, once the CBL was developed, it was easy to use. All the CBL values for the entire year were known in advance, which facilitated preparing and testing load management actions to have at the ready for when prices dictated undertaking them. Some participants employed sophisticated monitoring, analysis, and control technology to devise and carryout plans to respond to RTP prices. However, most devised simple rules taking advantage of the predetermined CBL values. They set a strike price for each day, and unless the RTP price passes through it, they go about business as usual. When the RTP price passes the strike price, an established load management response plan, using the predetermined CBL, was implemented. The simplicity and convenience of the predetermined CBL outweighed the inherent bias that arises from both random effects and from self-induced actions.

Recommendations for this Summer

The NYISO and the LSE Market Participants intend to implement the day-ahead bidding program this summer, beginning a soon as three months from now. Needless to say, simplicity and expediency should carry the day whenever practical. There are compelling reasons to adopt the 10-day simple averaging method used for the EDRP program for use in developing CBL. Despite its shortcomings, this dynamic CBL will minimize the effort required to subscribe customers to the program in the next few months. Moreover, its use in both programs will reduce the implementation and administration costs to LSEs, metering agents and the NYISO, especially with respect to using some other dynamic CBL formulation. Finally, bid floors or other measures can be used to mitigate the inherent bias of the dynamic CBL in some customer situations.

But, the dynamic CBL methodology imposes unacceptable bias on many customers who otherwise would be able to realize benefits from participation, and whose participation would generate large collateral benefits. Some LSEs may desire to undertake the time required to develop a predetermined CBL for those customers as part of their overall marketing campaign. Experience with RTP programs suggests that the investment in making sure the customer can participate and realize benefits will be worth the effort.

The NYISO should allow LSE's to use either the dynamic 10-day average CBL, as approved for the EDRP program, or a predetermined CBL for the period May through October 2001. The NYISO should promulgate rules and procedures to guide the development of the predetermined CBL, including, if practical, provisions for a simple, kW-based predetermined CBL. LSE's should be allowed to use alternative predetermined CBL methods that conform to basic guidelines set by the NYISO to promote diversity program offerings and advance the knowledge base needed to refine and improve the program.