A Probabilistic Method to Assess Power Supply Adequacy for the Pacific Northwest

John Fazio, Senior Systems Analyst

12/22/2011

A Probabilistic Method to Assess Power Supply Adequacy for the Pacific Northwest

Abstract 2
Introduction 3
What is Resource Adequacy? 4
Why the Pacific Northwest differs from most other Regions 5
What Conditions cause Problems for the Northwest? 5
Using Loss of Load Probability (LOLP) to Assess Adequacy 5
Not all LOLP Assessments are Equal 6
Calculation of LOLP for the Northwest 6
Defining Standby Resources 7
Assumptions about Non-firm and Market Resources 10
Determining an Appropriate Threshold for the LOLP 10
The Cost of Changing the LOLP Threshold 11
Translating LOLP into Deterministic Metrics 12
Annual vs. Seasonal Load/Resource Balance 15
Sustained-Peak Planning Margins 17
Other Probabilistic Methods 18
Loss of Load Probability 19
Loss of Load Expectation 20
Expected Unserved Energy 21
Other Metrics 22
Defining the Adequacy of a Resource Strategy 22
Appendix A: A Simple Example using Dice 26

ABSTRACT

This paper describes the methodology developed by the Northwest Resource Adequacy Forum to assess the adequacy of the regional power supply. This methodology was used to develop the Council's first adequacy standard (adopted in 2008).

Since 2008, the standard has been very useful to resource planners, although the assessments were somewhat cumbersome to calculate and results invariably led to misinterpretations. The analysis uses a probabilistic measure to assess the adequacy of the power supply. That probabilistic measure is then converted into more commonly used static measures of power supply, namely annual load/resource balance and hourly capacity planning margins. Comparison of these static measures to similar calculations reported in utility publications has always led planners into a quagmire of confusion, mostly because each set of calculations is used for different purposes.

Because of this and other issues, the Council chose to have this methodology reviewed by a peer group. Results indicated that while the 2008 standard was sound, it could be

improved. The resultant amended standard (adopted by the Council in 2011) is simpler and more useful than the 2008 version. Both the peer review report and the revised standard are available on the Resource Adequacy Forum web site (http://www.nwcouncil.org/energy/resource/Default.asp).

INTRODUCTION

Electricity does more than keep the lights on in the Pacific Northwest. It literally powers our economy. The absence or presence of an adequate electricity supply can either curtail or facilitate economic growth. In the worst extreme, an inadequate electricity supply can affect public health and safety, as in a blackout. Fortunately, such events are rare and when they do happen are most often caused by a disruption in the delivery of electricity (transmission lines), not the supply. However, there have been times – during extreme cold spells or heat waves – when the supply has been tenuous. The fact that most of the region's electricity comes from hydropower facilities presents unique challenges to power system planners and operators in the Northwest because periods of drought, which can severely limit hydropower production, are unpredictable.

There are a number of national, west-wide, regional and state efforts currently underway to define power system resource adequacy. The Energy Policy Act of 2005 mandates the Electric Reliability Organization (ERO), established by the Act to implement mandatory adequacy standards for the bulk-power system under the purview of the Federal Energy Regulatory Commission (FERC), "to conduct periodic assessments of the adequacy and adequacy of the bulk-power system in North America." The North American Electric Reliability Council (NERC), which was certified as the ERO on July 20, 2006, is in the process of developing a standard for resource adequacy assessments. FERC said in its final rule on implementation of the ERO provisions of the legislation that it intends to require the ERO to make recommendations where entities are found to have inadequate resources following the assessments.

In the West, the Western Electricity Coordinating Council (WECC) is developing guidelines to recommend appropriate methodologies for assessing resource adequacy. Although the NERC and WECC efforts act as drivers, momentum is also building within the region for a regional resource adequacy standard through the efforts of the Northwest Power and Conservation Council (Council) and the resurgence of utility Integrated Resource Plans (IRP).

This paper describes a probabilistic approach to assessing the adequacy of the Northwest's power supply. In the past, Northwest utilities have relied more on deterministic measures, such as annual average load/resource balance and capacity planning margins to assess supply adequacy. The use of probabilistic measures allows planners to better quantify the effects of future uncertainties, such as river flows, load variations due to temperature, variable wind generation and generator forced outages.

WHAT IS RESOURCE ADEQUACY?

Resource adequacy is not the same as power system adequacy, although it is a part of it. The terms *adequate* and *reliable* have specific meanings in the power industry. Adequacy is a component of adequacy. A power system is reliable if it is:

Adequate - the electric system can supply the aggregate electrical requirements of the customers at all times, taking into account scheduled and reasonably expected unscheduled outages of system elements.

Secure - the electric system can withstand sudden disturbances, such as electric short circuits or unanticipated loss of system elements.[1]

Adequacy refers to having sufficient resources – generation, efficiency and transmission – to serve loads. In determining adequacy, the Council uses sophisticated computer programs that simulate the operation of the power system over many different futures. Each future is simulated under a different set of unknown parameters, such as water supply, temperature, wind generation and thermal resource performance. The current Northwest standard calls for the power supply to have sufficient resources (both generating and conservation) to limit the likelihood of future years with significant curtailments to no more than five percent. The power supply must have sufficient capability to protect against both cold snaps in winter and heat waves in summer.

Security is achieved largely by having reserves that can be brought on line quickly in the event of a system disruption and through controls on the transmission system. These reserves can be in the form of generation or demand side curtailment that can take load off the system quickly. The North American Electric Reliability Corporation (NERC) and the Western Electricity Coordinating Council (WECC) establish reserve requirements. The reserve requirement is frequently expressed in terms of a percentage of load or largest single contingency, e.g., the loss of Energy Northwest's Columbia Generating Station. The reserves required for security are an additional resource requirement necessary for a reliable power system.

For the context of this paper, adequacy only refers to having an ample supply of generating resources and efficiency measures to "keep the lights on" a relatively high percentage of

[1] "Glossary of Terms," North American Electric Reliability Corporation, Glossary of Terms Task Force, August 1996 and also see http://www.nerc.com/page.php?cid=1%7C15%7C122.

time. It leaves out of the equation, the effects of bulk transmission and distribution systems, which typically cause most customer interruptions.

Why the Pacific Northwest differs from most other Regions

The Pacific Northwest power supply is mostly comprised of hydroelectric generation, which contributes anywhere from half to most of the region's electricity production on an annual basis. The variation in hydroelectric generation is due to two factors; 1) variability in precipitation, snow pack and river flow volume and 2) limited storage capability of the reservoir system. The combined storage capability of the US and Canadian reservoirs is only about 30 percent of the average annual runoff volume for the Columbia River. The aggregate peaking capability of the power supply, however, is much greater than the highest single hour load. Thus, the Northwest can be characterized as being capacity rich but energy limited.[2] Most other regions around the world are energy rich (assuming a controllable amount of fuel) but capacity limited (aggregate peaking capability is closer to the highest single hour demand).

What Conditions cause Problems for the Northwest?

There are a great number of events that could cause a stoppage of service. Because the Northwest is generally capacity rich, forced outages to generating resources are not the most dominant contributor to curtailments. Examining historical periods in the Northwest indicates that, in general, a combination of extreme temperature and poor water conditions tend to stress the system the most. Resource outages contribute to the problem but by themselves generally do not cause stress to the system.

For the assessment of power supply adequacy in the Northwest, four factors or random variables are examined; 1) runoff volume, 2) temperature, 3) forced outages for generating resources and 4) variability in wind generation.

USING LOSS OF LOAD PROBABILITY (LOLP) TO ASSESS ADEQUACY

Loss-of-load-probability is often used as a metric to assess the adequacy of a power supply. This paper addresses some commonly asked questions regarding this measure. In particular, what is it and how is it calculated? Are all LOLP assessments the same? What value of LOLP represents an adequate supply? Can we link a probabilistic LOLP metric to a

[2] Recent analysis, as describe in the Northwest Power and Conservation Council's Sixth Power Plan indicates that summer peaking issues are looming on the horizon.

more transparent and more easily calculated deterministic metric? Let's begin with a definition of loss of load probability or LOLP.

The loss of load probability is most often defined as the likelihood (probability) that system demand will exceed the generating capacity during a given period. For capacity-limited power systems, the concern is whether the daily peak load will exceed the system generating capability. For these systems, adequacy assessments are often expressed as the expected number of days per year in which peak load exceeds generation capacity. The metric used to express that measure is often referred to as the loss of load expectation (LOLE) or the loss of load hours (LOLH). These two metrics are often used interchangeably, even though they are not the same (see the section *Other Probabilistic Metrics*).

For energy-limited systems (namely hydro based systems), planners are more concerned about the supply of fuel (water). The LOLP gives resource planners an indication of how often the power supply will be insufficient to meet all customers' needs. Generally, a load serving entity will plan future resource acquisitions to keep the LOLP at an "acceptable" level. It should be noted that this definition of LOLP does not include uncertainties surrounding the availability of the bulk transmission system. (A simple example using dice is provided in Appendix A).

Not all LOLP Assessments are Equal

LOLP assessments for different power systems are not necessarily the same. LOLP is sensitive to the uncertainties that are modeled in the simulation program. For example, in some capacity-limited regions, thermal resource performance (forced outage rate) is the only uncertainty incorporated into the assessment. For the Northwest, three other uncertainties are included, namely natural river flows, temperature and wind generation. In this example, the LOLP calculation for the Northwest cannot be compared directly to one calculated for the hypothetical capacity-limited region described above. In order to be able to compare these two assessments, the capacity-limited region must also include temperature and wind generation uncertainty in its calculation.

The same argument can be made about LOLE, LOLH and other adequacy metrics. If the intent is to measure the likelihood of possible shortfalls, then *all relevant* uncertainties must be included in the adequacy metric calculation. If that is done, then adequacy measures can be compared across different regions.

Calculation of LOLP for the Northwest

LOLP is generally calculated by means of stochastic simulation models. These computer models simulate the operation of a power system over many potential future conditions and record all hours when demand cannot be met. The ability of a simulation model to produce realistic results depends on how many relevant parameters are simulated, how well they are simulated and how good the associated data is. Genesys, the model used to analyze the Northwest power supply, simulates the operation of thermal and hydroelectric resources used to meet regional demand.

Genesys first performs a monthly dispatch to assess how much hydro generation is available for that month. Once the monthly hydro generation is established, the model drops into an hourly simulation where individual thermal plants are dispatched along with the *aggregate* hydroelectric system (Genesys does not simulate the operation of individual hydroelectric projects in the hourly mode). The available monthly hydro generation is shaped to meet hourly loads as best it can without violating its minimum and maximum hourly generation limits. Hourly hydro generation limits are derived from a separate hourly hydro model that simulates the operation of individual projects and takes into account various operational constraints and reserve requirements. Some federal hydro projects, for example, carry within-hour balancing reserves for wind generation.

Any hour in which load cannot be served is recorded as a shortfall. However, since Genesys does not model standby resources (defined in more detail below), resulting shortfall events are first screened to eliminate those that could be avoided by dispatching standby resources. The LOLP is calculated as the number of simulated years, which have at least one shortfall event that exceeds the aggregate standby resource capability, divided by the total number of simulated years. For the Northwest, a five percent value has been assigned as the adequacy threshold (i.e. LOLP must be five percent or less for the power supply to be adequate).

Defining Standby Resources

Standby resources are generating resources and demand-side management actions, contractually available to Northwest utilities, which can be accessed quickly, if needed, during periods of stress. These resources are intended to be used infrequently and are only used as a last resort. Ideally, standby resources can be identified and their aggregate energy and capacity capability can be calculated. However, because some of this information is difficult to acquire, the Council (since 2001) has assumed that the aggregate capability of standby resources is 28,800 megawatt-hours (1,200 megawatts over 24 hours) of energy over a summer or winter period and 3,000 megawatts over any hour.[3]

[3] The Resource Adequacy Forum is in the process of identifying standby resources and calculating their aggregate energy and capacity capability.

Table 1 and Figure 1 below identify the relationship between the *energy* capability of the standby resources and the resultant LOLP. It should be noted that only one study case was used to generate the graph in Figure 1 and the data in Table 1. This study case has an approximate annual load/resource deficit of 1,200 average megawatts and, using the currently defined threshold, about a 4 percent LOLP. This case is considered to be adequate under the 2008 standard. It is interesting to note how widely the LOLP changes depending on the size of the standby resource. (The term "seasonal curtailment" in Figure 1 should be interpreted as the energy capability of standby resources.)

Table 1
Energy Capability of Standby Resources vs. LOLP

Maximum Energy (megawatt-hours)	LOLP (%)
0	36
4,000	18
10,000	12
28,800	4

Figure 1
Relationship between Energy Capability of Standby Resources and LOLP

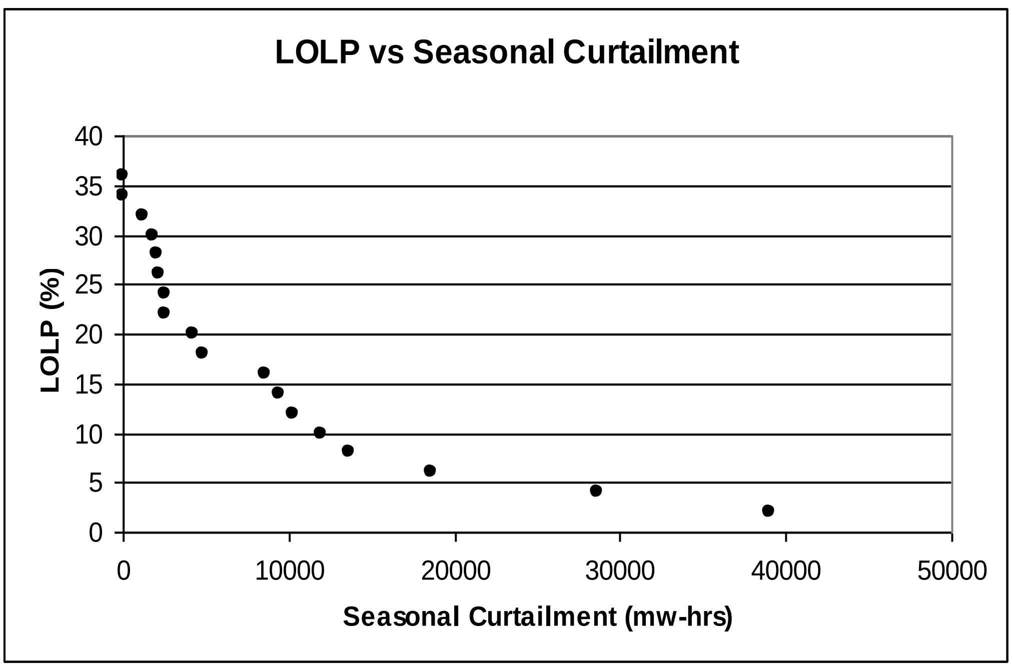

Assumptions about Non-firm and Market Resources

As part of an adequacy assessment, all available resources should be considered. This includes market resources and "non-firm" resources (in the Northwest this means hydro energy above the critical year generation). The question being addressed by this analysis is, "can we meet all load requirements, cost notwithstanding?" Thus, this analysis is not intended to produce a resource planning target because it does not include economic considerations. To calculate the regional LOLP, the Council first determines the amount of market supply (both in region and out of region) that is available by month, regardless of cost.

Relying solely on market resources will likely keep the average cost of the power supply low but will have a higher likelihood of high cost years. Not relying on the market at all will yield the least risk with respect to high cost years but will also result in the highest average cost. An adequacy assessment, as defined by the Resource Adequacy Forum, does not take these economic factors into account. However, this method can be used to aid in the development of integrated resource plans by first defining (policy call) how much market supply a utility is willing to rely on, then assessing the LOLP. If the LOLP is higher than five percent, additional resources can be added to bring the system back into acceptable limits.

DETERMINING AN APPROPRIATE THRESHOLD FOR THE LOLP

The question of where the five percent threshold for the LOLP originated is often asked but never properly answered. It seems to date back to the early half of the 20th century when planners deemed it appropriate to limit their tolerance to unwanted curtailments to only once in 20 years. This once in 20 year threshold is converted into the five percent limit for LOLP. However, there remains ambiguity regarding the once in 20 year tolerance level. Does that imply that we will only tolerate one "curtailment event" during a 20 year period or does it mean that we can tolerate one bad year (perhaps with multiple curtailments) in a 20 year span? The Council has chosen the latter definition, which is not consistent with the majority of regions in the United States that use a one-day-in-10 year threshold. However, ambiguity also surrounds the one-day-in-10 year limit. Does it literally mean one day of curtailment (probably not) or one event per 10 year period (probably yes)? (See the section *Other Probabilistic Metrics* for more discussion on this.)

For the Northwest, the Council has decided that our tolerance for unwanted events is limited to one "bad" year during a twenty year span. This definition makes a lot of sense for the Northwest because the most likely cause of shortfalls is a shortage of water. A dry season can last from December through the end of summer and possibly beyond. Thus, during a dry year, the region will likely experience multiple shortfall events. This is the primary reason that the Council adopted a five percent maximum limit for LOLP as defined by the number of "bad" years divided by the total number of years simulated.

The Cost of Changing the LOLP Threshold

One way to reduce the LOLP of an inadequate power system is to add new resources. An alternative, of course, is to reduce demand. But for this analysis, a hypothetical inadequate system was analyzed. The base case LOLP is 21 percent, well above the currently used five-percent target. New resources were added incrementally to reduce the LOLP. In this case combined cycle combustion turbines were used. An equal amount of turbine capacity was added to both the east and west side of the Northwest region. The turbines were added in 100-megawatt increments for both parts of the region.

Figure 2 below illustrates how the LOLP is reduced with the addition of new resource capacity. Clearly the relationship is not linear, but looking at the total change, it took 2,000 megawatts of capacity to reduce the LOLP by 16 percent. On average, each 125 megawatts of additional capacity reduces the LOLP by one percent. Assuming that the levelized life-cycle cost of a turbine is about $55 per megawatt-hour (based on a 70 percent capacity factor) the cost of reducing the LOLP from 21 to 5 percent is about $670 million per year. On average, the cost of reducing the LOLP by one percent is about $42 million per year. Obviously, a more sophisticated and detailed analysis should be done to make a more precise assessment of the cost to lower LOLP. (It should be noted that these calculations were done a number of years ago and have not been updated. Thus, this is only an illustration of the type of analysis that can be done with Genesys.)

Figure 2

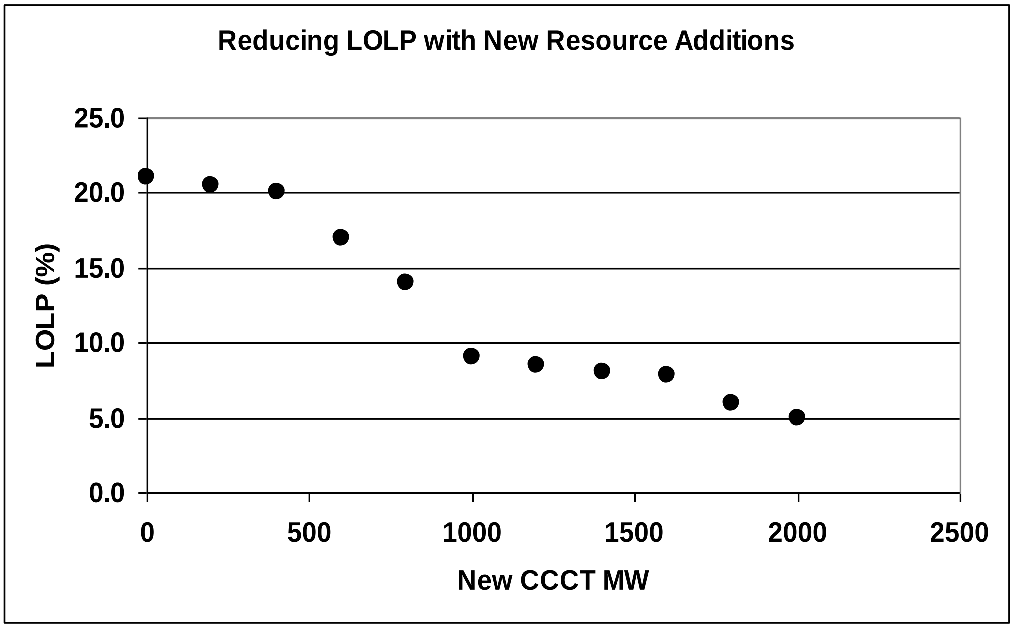

TRANSLATING LOLP INTO DETERMINISTIC METRICS

A link can be established between the LOLP and the deterministic load-resource balance energy metric. This is a similar process used in other regions where a probabilistic adequacy metric is translated into a deterministic capacity planning margin (a capacity metric). For the energy-constrained Northwest, the five percent LOLP can be translated into a load-resource-balance adequacy threshold.

Of course, many parameters will affect this conversion, in particular, assumptions about the market supply plays a very important role in assessing the LOLP.[4] The relationship among these three variables (load-resource balance, market supply and the LOLP) is shown in Figure 3. The x-axis represents the available supply of out-of-region *capacity* over the winter period. The y-axis represents the LOLP. Each curve in this figure shows the relationship between available out-of-region supply and the resultant LOLP for a *constant load-resource balance*.

For example, the third curve from the left in Figure 3 represents a case in which the load-resource balance deficit is about 2,000 average megawatts. Extrapolating this curve up to the point where it crosses the vertical axis yields an LOLP of approximately 30 percent

[4] Interruptible loads and demand-side management programs must also be taken into account.

(labeled P1) – well above the five percent threshold. This means that with no available out-of-region winter capacity, planning to a deficit of 2,000 average megawatts leaves the region extremely inadequate.

Assuming that 3,000 megawatts[5] of out-of-region winter surplus capacity is available (labeled P2), the LOLP drops to about 12 percent – much closer to the desired threshold. Assuming 4,000 megawatts of available out-of-region supply (labeled P3) brings the LOLP to about a five percent value. Thus by knowing how much out-of-region surplus winter capacity is available; we can determine the proper load-resource balance to plan to, which will yield a five percent LOLP for the Northwest.

Figure 4 illustrates the relationship between out-of-region market supply and load-resource balance for a constant five percent LOLP. Any point along the curve represents, by definition, an adequate Northwest supply. Thus, if 3,000 megawatts of out-of-region supply were available, the Northwest should plan to a load-resource balance deficit of about 1,500 average megawatts in order to achieve a five percent LOLP. This well-defined relationship allows us to use the annual load/resource balance metric and associated threshold with confidence because we know that it is supported by the more sophisticated LOLP methodology.

Because it would be difficult for individual utilities to plan to a *regional* load-resource balance threshold under critical hydro conditions, an alternative would be to plan to a zero load-resource balance but assuming better than critical water conditions. Figure 5 shows the relationship between out-of-region market and adverse water conditions for a constant five percent LOLP. The 100 percent adverse water condition is defined to be the critical hydro condition. Average water conditions appear in Figure 5 at the 50 percent mark on the x-axis. Thus, if 3,000 megawatts of out-of-region capacity were available, the region should plan to approximately the 85th percentile adverse water condition to achieve a five percent LOLP.

[5] This does not mean that the region would import 3,000 megawatts each hour of the winter period. It simply means that 3,000 megawatts of surplus out-of-region capacity is available. During some hours, the entire amount may be imported but on average, over all hours and all simulated futures, only about [1,200] average megawatts of energy is imported.

Figure 3

Relationship among LOLP, Load-Resource Balance and Out-of-region Supply

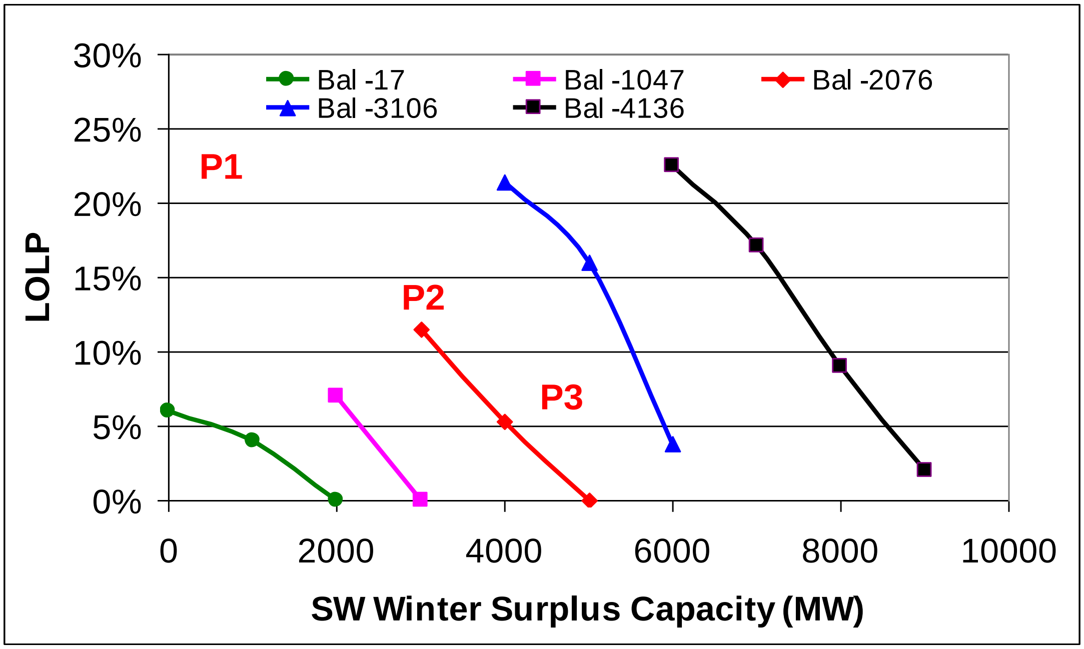

Figure 4

Relationship between Out-of-region Supply and Load-Resource Balance

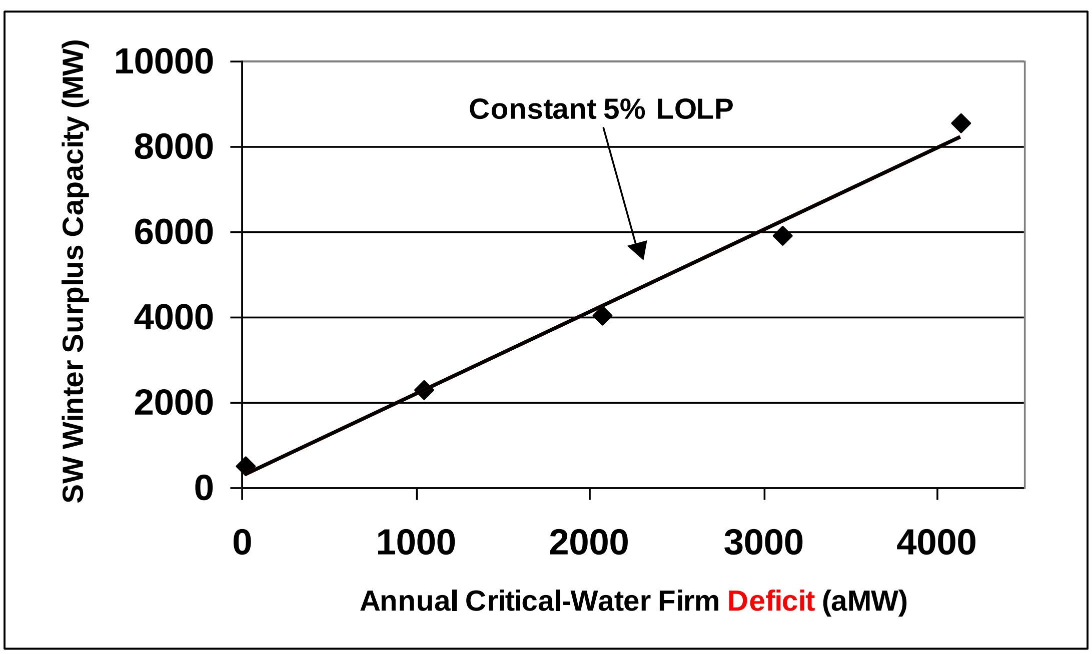

Figure 5
Relationship between Out-of-region Supply and Adverse Water Condition

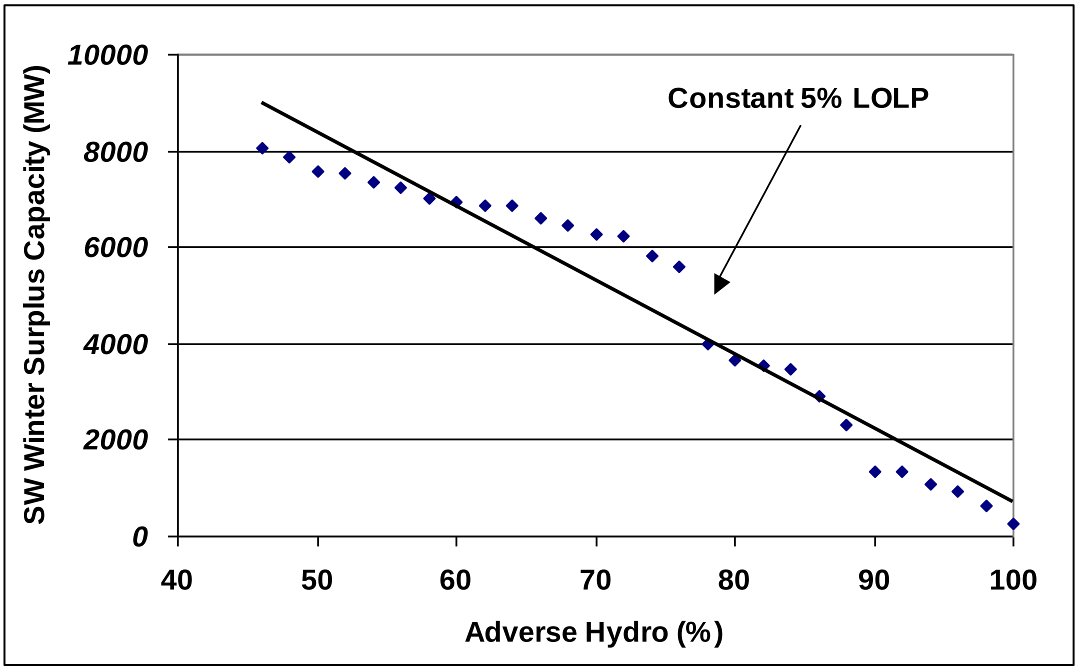

Annual vs. Seasonal Load/Resource Balance

Should the Northwest use a seasonal load/resource balance as its deterministic energy adequacy metric? This seems appropriate since the Northwest is a winter peaking region (at least for the time being) and adequacy assessments should focus on the peak demand period. The problem with using a winter season load/resource balance, however, is that it requires more detailed information, some of which may be hard to get. For example, maintenance schedules would have to be known more precisely. For an annual load/resource balance calculation, on the other hand, only the duration of maintenance needs to be known, not the schedule.

Other problems related to a seasonal load/resource balance calculation include de-rating thermal resource output to capture forced outage rates, reassessing generating capability as a function of temperature and applying other seasonal operating constraints (such as air quality or noise). Using an annual load/resource balance minimizes the relative effect of these approximations that must be made.

One reason for using an annual load/resource balance is that it is well understood and has been calculated in the region for over 30 years. Remember that we are using the load/resource balance merely as a surrogate for the more precise LOLP metric. By running many simulations involving many different future resource and load scenarios, a relationship can be established between the LOLP and the annual load/resource balance (as discussed in the previous section). A similar relationship could be derived between the LOLP and a seasonal load/resource balance. In the end, however, it doesn't matter whether an annual or seasonal number is used because the two are related in a very linear fashion. Figure 6 below shows the relationship between the annual and seasonal load/resource balance for a number of cases. It should be clear from that figure that the relationship between these two variables is very linear, meaning that we can link either the annual or the seasonal value to the LOLP.

The only caution we need to concern ourselves with is that this linear relationship depends on the demand shape and resource generation shape in the Northwest. As the demand shape changes, say air conditioner penetration rates greatly increase over time for example; the slope of the line in Figure 6 might change (but there is no evidence that the relationship will become non-linear). Also, should the constraints on the hydro system change drastically, the relationship between annual and seasonal load/resource balance may also change, but that can be quantified. If these events were to occur, reassessing the relationship between LOLP and the load/resource balance of choice allows the region to adjust the adequacy threshold.

An argument for using a seasonal load/resource balance as opposed to an annual balance is the idea of perception. Some say, rightly so, that when the region is surplus on an annual basis, it could very well be deficit in the winter months. Figure 7 illustrates this for a case in which the annual load/resource balance is about 1,000 average megawatts but the winter months are deficit. However, this picture doesn't necessarily help with perception. If the region were 1,000 average megawatts surplus on an annual basis, the deficits in the winter months should be of no concern partially because the monthly load/resource balances are calculated using only the critical hydro condition. Out-of-region market supply should almost always be available because of the diverse load patterns between the Northwest and the Southwest. Also, some maintenance can often be deferred during emergencies, thus also alleviating potential shortfalls. In this particular case (annual load/resource surplus of 1,000 average megawatts), the region has no need to acquire new resources. Yet by looking at the monthly balances for winter in Figure 7, that conclusion is not apparent. In fact, using an annual load/resource balance is not necessarily intuitive. From previous analysis, it appears that the region can plan to an annual load/resource *deficit* of about 1,500 average megawatts to maintain an adequate supply. This statement does not appear to be intuitive because the computation of the load/resource balance does not include any contribution from available out-of-region market supplies. So, for the sake of transparency and ease of use, the annual load/resource balance should be used as the energy metric. (Understanding, of course, that the link between the LOLP and the annual load/resource balance must be re-evaluated on a least a yearly basis).

Figure 6

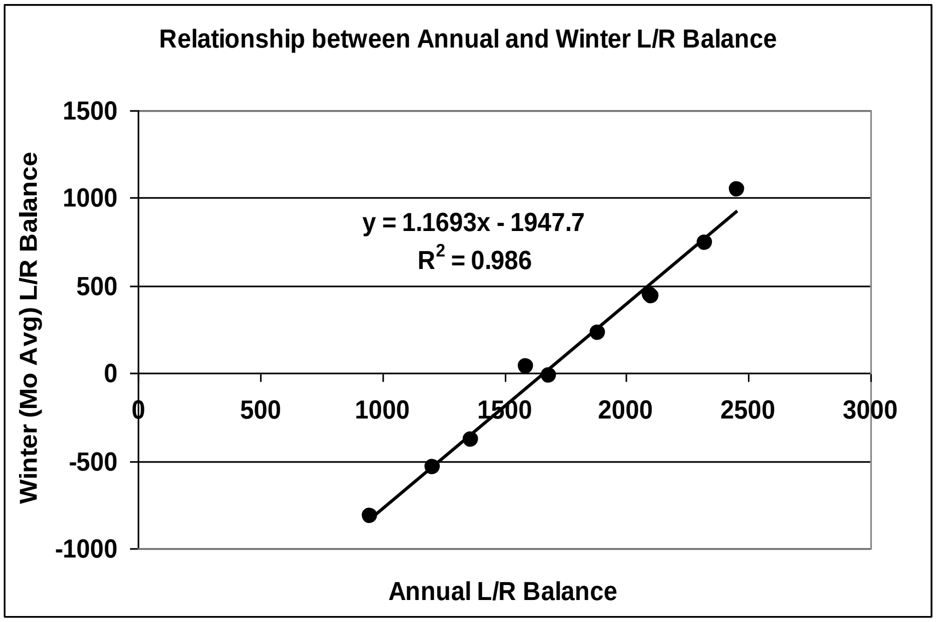

Figure 7

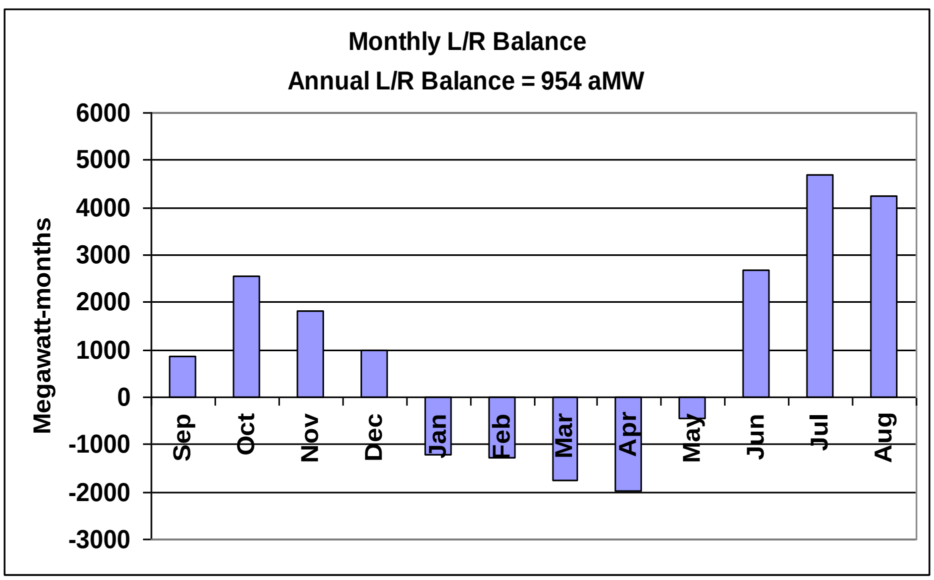

Sustained-Peak Planning Margins

A similar relationship between the five percent LOLP and capacity planning margins can be made. Capacity planning margins are defined as the amount of surplus generating capability over a single hour (or over multiple hours, which is referred to as a sustained peak period). For the Northwest, given the limited amount of reservoir storage and hydroelectric constraints, an 18-hour sustained-peak period has been determined to be the most appropriate for capacity assessments. This 18 hour period spans the 6 highest load hours over three consecutive days.

As with the relationship between the LOLP and annual load/resource balance, the resulting capacity margin thresholds will be a function of the assumed market supply. Once that assumption is made, sustained-period capacity planning margins can be calculated. Using this technique and based on the 2008 standard, the translation of the five percent LOLP threshold yields a 23 percent and a 24 percent sustained-period planning margins for winter and summer, respectively.

OTHER PROBABILISTIC METHODS

The main aspects of adequacy include the frequency, duration, and magnitude of interruptions to electricity service. There are many approaches and metrics for evaluating system adequacy throughout the utility industry. The nation-wide trend appears to be the use of probabilistic metrics, which typically use Monte Carlo simulation techniques to simulate future system performance under various uncertain variables (such as load variations and thermal resource availability). Four common probabilistic metrics are described below. The basic concepts and formulas are provided for each, although many variations are being used by planners.

Loss of Load Probability (LOLP), in units of *percent*, measures the probability that at least one shortfall event will occur over the time period being evaluated. By definition, since a probability must be greater than or equal to zero and less than or equal to one, LOLP is calculated as the number of simulations in which a shortfall occurs divided by the total number of simulations. It does not reflect the frequency of events because simulations with one or multiple shortfall occurrences are counted equally. LOLP also provides no information regarding duration or magnitude of resource shortfalls.

Loss of Load Expectation (LOLE), in units of *days per year*, is calculated as the number of days in which a shortfall occurs over every simulation divided by the total number of years simulated. Historically, many utilities have used a one-day-in-ten year threshold (or 0.1 day/year) to plan for adequacy. This, however, can be misleading because multiple shortfall *events* can occur during a single day and a single *event* can last longer than one day. Originally, before the advent of the fast computers we use today, only the peak hour of each day was examined, thus equating an *event* to a *day*. However, most utilities now simulate the operation over each hour of the year and use the LOLH metric (described below). The LOLE provides no information regarding duration or magnitude of resource shortfalls.

Loss of Load Hours (LOLH), in units of *hours per year*, is calculated as the number of hours in which a shortfall occurs over every simulation divided by the total number of years simulated. Historically, many utilities have translated the one-day-in-ten-year threshold into a 0.1 day/year or into a 2.4 hours/year threshold for LOLH. As noted above, this translation is not valid since a typical shortfall event does not last 24 hours. A more typical duration for a shortfall event is on the order of 8 hours. Thus, if the intent is to limit shortfall *events* to one in 10 years (or 0.1 per year), the correct LOLH threshold is 8 hours/10-years or 0.8 hours/year. In this sense, the LOLH is a more precise metric for assessing adequacy than the LOLE. Like the LOLE, the LOLH provides no information regarding duration or magnitude of shortfalls.

Expected Unserved Energy (EUE), in units of *megawatt-hours*, measures the expected amount of energy (in megawatt-hours) not being served per year (or per hour). It is calculated by adding up all of the unserved energy over every simulation and dividing by the total number of years simulated (or by the total number of hours simulated). EUE provides some indication of the magnitude of shortfalls but only in aggregate. It does not reflect the frequency, duration or magnitude of individual shortfall events.

Overall, since no single metric provides all meaningful information, consideration should be given to all aspects of adequacy that planners value. The following provides details on how these adequacy metrics are calculated[6].

Loss of Load Probability

Loss of Load Probability (LOLP) is calculated by dividing the number of simulations with at least one shortfall event by the total number of simulations. If each simulation represents a potential "future" for customers, then LOLP is the probability that customers will face an interruption at some point in the future. LOLP does not represent the probability of a single shortfall event occurring. The analysis is typically performed with an hourly level of granularity over an entire year, that is, a single year Monte Carlo study with 3,000 simulations means that the year's operation is simulated 3,000 times for each of the 8,760 hours in the year. In some cases a utility or region may only want to focus on its most critical part of the year, such as winter for the Pacific Northwest. In this case, hourly simulations only need to be done over the winter months and the resulting LOLP is thus a winter adequacy measure only.

What constitutes an event must also be defined and can vary among utilities and regions of the country. Planners usually include standby resources in their analyses, that is, resources

[6] The concepts illustrated here are based on "Use of Monte Carlo Simulation In Teaching Generating Capacity Adequacy Assessment," by R Billinton and L. Gan, published in Transactions on Power Systems, Vol. 6, No. 4, November 1991.

(including demand cutback agreements) that can be called upon during emergency situations. Alternatively, if standby resources are not modeled, shortfall events identified in a simulation may be screened to sort out those events that would be resolved by calling on standby resources.

Another element to assessing adequacy is whether consecutive hours are treated as individual events or as a single event. For example, a continuous 8 hour shortfall could be considered as 8 events, or just a single event, or not an event at all if the minimum threshold is 24 hours. Understanding the definition of "event" is important to understanding differences or similarities in various LOLP studies. The following equation is a generalization of how LOLP is calculated:

$$LOLP = \frac{\sum_{i=1}^{N} S_e}{N}$$

Where:

LOLP = Loss of Load Probability (%)

S_e = a simulation in which at least one significant event occurs. A significant event occurs when load and operating reserve obligations exceed resources including standby operations (or event threshold limits).

N = the number of Monte Carlo simulations for the period, which is typically one year using hourly level of granularity

Loss of Load Expectation

Loss of Load Expectation (LOLE) is calculated as the number of days in which a shortfall occurs divided by the total number of years simulated. This metric indicates the frequency of occurrence for days with shortfall events. As mentioned above, however, LOLE is not a frequency measure for individual events because multiple events can occur during a single day and a single event can last longer than one day. Historically, utilities have used a "one day in ten year" threshold for adequacy planning. LOLE provides a different measure of adequacy than the LOLP metric but neither provides any indication of duration or magnitude. As with the LOLP metric, LOLE can be limited to counting only "significant" shortfall events, that is, those that exceed minimum capacity, energy and/or duration thresholds.

A variation of LOLE is LOLH, which is defined as the loss of load hours. LOLH (in units of hours per year) is calculated by dividing the total number of hours in which a shortfall occurs by the total number of years simulated. This metric can be (and sometimes is) converted into a probability metric by dividing the resulting number of hours per year by 8,760 hours to yield the likelihood of experiencing a shortfall in any given hour. For a system that satisfies a "one in ten year" adequacy threshold, the corresponding LOLH value is often converted to 24 hours per 10 years or 2.4 hours per year. However, as described earlier, this translation is not correct and should be more like 0.8 hours per year.

LOLE is calculated as follows:

$$\text{LOLE} = \frac{\sum_{i=1}^{N}\sum_{y=1}^{Y}\sum_{d=1}^{D} D_e}{\text{Ny}}$$

Where:

LOLE = Loss of Load Expectation (days/year)

De = a day in which at least one significant event occurs (based on the definition of a shortfall event)

N = the number of Monte Carlo simulations for the period, which is typically one year using hourly level of granularity
Y = number of years in the study
Ny = the total number of years simulated in the Monte Carlo study
D = the number of days in each year that are simulated

Expected Unserved Energy[7]

LOLP and LOLE are metrics more closely associated with the frequency aspect of adequacy. Neither provides any indication of the size or magnitude of potential shortfalls nor of the duration of shortfalls. Expected Unserved Energy (EUE) is defined to provide some measure of the magnitude of shortfall events. It is calculated as the sum of all unserved energy (in megawatt-hours) over all hours of the simulation divided by the total number of hours simulated. The resulting expected megawatt-hour loss per hour value is sometimes

[7] Also see "Justification for a NERC Resource Adequacy Assessment Model, A NERC Staff White Paper," Bob Cummings, Mark Lauby, John Seelke, February 28, 2007, Revised July 31, 2007 pg 3 http://ewh.ieee.org/cmte/pes/rrpa/RRPA_files/2.28.07%20rev%2007.31.07%20Justification%20for%20a%20NERC%20Resource%20Adequacy%20Model.pdf.

translated into an expected megawatt-hour loss per year (for example, a five MW-hour/hour expected loss becomes five MW-hour/hour x 8760 hours/year or 43,800 MW-hours/year of expected energy loss). EUE provides no indication of event frequency or duration. EUE is calculated as follows:

$$\mathrm{EUE} = \frac{\sum_{i=1}^{N}\sum_{y=1}^{Y}\sum_{d=1}^{D}\sum_{h=1}^{H} E_h}{\mathrm{Nh}}$$

Where:

EUE = Expected Unserved Energy (MW-hours/hour)
N = the number of Monte Carlo simulations for the period, which is typically one year using hourly level of granularity
Y = number of years in the study
D = number of days in each year that are simulated
H = number of hours in each day that are simulated
Eh = the amount of unserved energy for this hour (in megawatt-hours)

Nh = the total number of hours simulated in the Monte Carlo study.

Other Metrics

Numerous other metrics can be defined to address different aspects of adequacy. As stated earlier, the three key parameters of shortfall events are frequency, duration and magnitude. Any number of adequacy metrics can be defined using some combination of these three key parameters. The three adequacy metrics described above are the most commonly used metrics today but even in aggregate they do not fully describe the characteristics of shortfall events. Each utility or region must decide what factors are most important for its own adequacy planning. For example, most utilities are capacity (or machine) limited and thus will use metrics that relate more to hourly needs. Energy-limited regions, such as the Pacific Northwest, may choose a different metric that better addresses annual energy needs.

DEFINING THE ADEQUACY OF A RESOURCE STRATEGY

The Resource Adequacy Forum has developed a method to assess the adequacy of the Northwest's power supply. More specifically, it has defined a probabilistic measure to gauge whether Northwest's resources will sufficiently satisfy the region's needs. It focuses only on the adequacy of electricity supply and does not take transmission outages into account (it does capture variation in transmission capacity for the east-to-west regional interties).

The Northwest Resource Adequacy Standard uses the probabilistic measure defined by the Forum to assess whether existing resources will be sufficient through the next five year period. That assessment only takes into account existing resources and new resources that are expected to be completed and operational during that time period. If the power supply is deemed to be inadequate (e.g. LOLP greater than five percent), then specific actions are initiated. Those actions include reporting the known problem, validating load and resource data and identifying potential solutions.

The process described above is intended to be an early warning for the region that indicates when resource development does not sufficiently keep up with demand. Although similar, the assessment of a resource strategy differs in significant ways. First, a resource strategy spans a much longer time period, namely 20 years for the Council's power plan. Second, a strategy implies that resource development will be dynamic, in other words, it does not identify specific resources and specific build dates. Rather, the strategy identifies a supply of cost-effective resources that can be acquired as future conditions warrant. One can extract a single resource plan out of a particular resource strategy and then assess the adequacy of that single plan but that is not the same as assessing the adequacy of the strategy itself.

The adequacy measure, as adopted by the Council, assesses the sufficiency of a specific set of resources combined with a specific forecasted demand by simulating the operation of those resources over many different futures. In those futures, water conditions, temperatures (which affect load), wind generation and thermal resource availability can vary. Based on those random variables only, a loss of load probability is calculated and compared to the five percent maximum allowed under the standard.

The five percent LOLP threshold can be translated into deterministic metrics, which are more easily used for assessing adequacy or for incorporation into resource planning models. For example, a power system will provide an adequate supply of energy for the region when the average generation of existing resources plus about 1,500 average megawatts of market supply equals the average annual load. Similarly, the system will provide an adequate supply of peaking capability when the surplus sustained-peak generating capability is 23 percent in winter and 24 percent in summer. In each of these cases, the resulting LOLP will be five percent.

These calibrated deterministic metric thresholds are easily incorporated into the Council's Resource Portfolio Model (RPM). That model simulates a wider variety of future conditions with many more future unknowns than the Genesys model, which is used to assess LOLP.

The RPM acquires resources based on economic considerations but if those resources do not measure up to the deterministic adequacy metric threshold, the model will add resources until that condition is satisfied. In this way, the Council can be sure that each resource plan examined under any particular resource strategy will be adequate, at least for energy supply.

The problem is that currently, the only adequacy metric incorporated into the RPM is the energy measure. It may be possible (but unlikely) that some resource plans generated by the RPM may not meet the peaking adequacy thresholds. Given that the region is transitioning from a winter energy-limited system into a summer capacity-limited system, this omission from the RPM needs to be addressed. Perhaps future versions of the model can also include measures to test the peaking adequacy of various resource plans.

However, the question at hand is how to assess the adequacy of a resource strategy developed by the RPM. One suggestion is to assess the adequacy of each resource plan (types of resources and build dates) for all 750 simulated futures for each strategy, using the deterministic adequacy metrics. This is time consuming but very doable. However, it is not clear what an acceptable result would be. Do all 750 plans need to be adequate or would it be acceptable if only 95 percent of them were adequate? Another option would be to assess the adequacy of the "average" build-out schedule of the strategy. If this average scenario is adequate does that imply that the strategy is adequate?

The problem is that the RPM simulates future conditions with many more random variables than does Genesys. The most important variable, from an adequacy point of view, is probably long-term load uncertainty. This is not the uncertainty in demand caused by variation in temperature but rather the potentially much larger change in demand due to economic or other factors. A result of this is that the RPM will simulate situations when the region will under or overbuild, much like it has in real life over the past 50 years. There really is no way to avoid such conditions because we cannot accurately forecast all future conditions, especially demand. We could have the RPM calculate a loss of load probability for its all of its plans in each strategy but that calculation could be misleading. Although labeled LOLP, the RPM version provides a vastly different measure of the power supply than does the Genesys LOLP because the random variables are different. At this time it is unknown what a reasonable RPM LOLP value would be or whether it would ever be meaningful, since the RPM is not an hourly simulation model and thus can only approximate peaking operations.

So what do we do?

The first thing to remember is that the real scope of this power plan has a five year time period. We will revisit these questions five years from now. So, potential inadequacies in the later years of the study horizon may be interesting but are unlikely to change the five-year action plan. Thus, if we are to assess the adequacy of all (or some) of the resource build-outs from an RPM strategy, we should only focus on the first five years. It seems to me that most of those plans should pass the Genesys adequacy test (at least the deterministic ones). If a significant number of those plans fail the test, we should ask ourselves why that is and perhaps change our five-year action plan to address the problem. However, this does not appear to be the case for the draft power plan.

There is the additional issue of whether the adequacy of all RPM strategies should be assessed. Some of those strategies are based on assumptions that have little or no likelihood of being realized. In those cases, it makes no sense to spend the time calculating adequacy, especially because they do not drive the action items in the plan. Thus the current recommendation is to simply use some form of deterministic adequacy metric inside of the RPM to dynamically test for adequacy during the analysis. It is not recommended that we take any of the "build-out" cases out of the RPM and assess the LOLP specifically.

APPENDIX A: A SIMPLE EXAMPLE USING DICE

If temperature were the only random variable, then utilities could plan to have sufficient resources to cover a 1-in-20 year temperature deviation. If forced outages were the only random variable, then they could plan to cover a 1-in-20 year outage event. And, if water were the only random variable, utilities could base their resource acquisitions on a 1-in-20 year adverse water condition.

But, unfortunately, these three random variables all need to be considered in combination. We can assume for now that they are independent variables, that is, the likelihood of one does not affect the likelihood of the others.

Perhaps a good way to start is to take a very simple example. Let's use a pair of dice. We know that there are 36 potential combinations, which yield 11 distinctive outcomes (totals from 2 to 12). The likelihood of getting a 2 is 1/36 because there is only one combination out of 36 that yields that result. The same is true for 12. The likelihood of getting a 7 is 1/6 because there are 6 different combinations that yield 7 (1-6, 6-1, 2-5, 5-2, 3-4, 4-3).

Let's say that a "bad" outcome is when we get a 1 on either die. The potential combinations include 1-1, 1-2, 1-3, 1-4, 1-5, 1-6, 2-1, 3-1, 4-1, 5-1, and 6-1. That amounts to 11 combinations out of 36 possibilities, so the probability of a bad outcome is 11/36 or about 31 percent (assuming we do not have weighted dice). The likelihood of getting a 1 on one die is 1/6, the likelihood of getting a 1 on both dice is 1/36 and the likelihood of getting a 1 on at least one die is 11/36.

Now let's say that one die represents the random variable for temperature and the other represents hydro conditions. Planning to cover a 1-in-6 temperature event coincidentally with a 1-in-6 water event is effectively planning to cover a 1-in-36 event. If we want to cover any combination that produces either a bad temperature or a bad water event then we should plan to protect against an 11/36 likelihood event. So, how do we do that?

Let's assume that each condition for temperature and water affect the load/generation by the same amount. If that is true, then there is a minimum change that we are trying to protect against and that is the effect of getting a 1 for either the temperature die or for the water die. The "reserve margin" needed to protect against getting a 1 for either temperature or water is well defined. The problem is that there may be other combinations of temperature and water events that also yield the same change to load, generation or both.

Knowing the change in load/generation that we wish to protect against allows us to set the reserve margin but it doesn't help us decide what particular temperature or water event we are protecting against. For example, we already know that protecting against a 1-in-20 year water event coincidentally with a 1-in-20 year temperature event is too conservative. If our goal is to protect against a combined 1-in-20 year event for water and temperature, we could say that we are protecting against a 1-in-4.5 year temperature coincidentally with a 1-in-4.5 year water event (1/4.5 times 1/4.5 is equal to 1/20). But that doesn't count the years when we might have a 1-in-3 water year and a 1-in-7 temperature year, which might also cause a problem. The hope is that if we choose the "right" combination of water and temperature conditions, we may be able to assess the "right" level of protection -- that is, one that is consistent with a five percent loss-of-load probability.

Let's try a slightly different approach. Let's assume that we can enumerate all the possibilities (like with the dice example). Then for each combination, we would calculate the required magnitude of resource to cover the deviation. For our dice example, we would have 36 possibilities, each one with a potentially different resource size requirement. If we sort the 36 magnitudes from highest to lowest, we can easily draw a line where the 95th percentile is -- that is, we want the 95th percent highest magnitude, which tells us how much resource we need to cover 95 percent of the possible future contingencies. The problem with this method, unfortunately, is that it is difficult if not impossible to enumerate every possible contingency.

There may be a solution, however. There is a good deal of literature on this subject dating as far back as the 1960s. The idea is to use a booth-baleriaux8 method to convolve probabilities associated with uncertain loads and with hydro and thermal availability.

We can determine a magnitude for the planning margin by assessing the load and/or generation deviations that we would like to protect against. If we choose carefully, we can be consistent with the five percent LOLP assessment.

We must remember that the LOLP analysis does not assess a planning margin by adding up components -- for example, X percent to cover a 1-in-W water event plus Y percent to cover a 1-in-T temperature event. The LOLP analysis is simply protecting against any combination of temperature and water events that yields a 1-in-20 year load/generation deviation.

[8] Links to sites that offer representative papers on this subject include;
http://www.ohiolink.edu/etd/view.cgi?ohiou1182181023,
http://www.dis.anl.gov/publications/articles/IEEE_EMCAS_Expansion_v5.pdf,
http://en.scientificcommons.org/732963