The Benefits of Goal-Based Planning in Times of Uncertainty

Pamela Lesh, Graceful Systems LLC

Pamela@gracefulsystems.com 503-701-2875

Abstract: The predominant mode of resource planning by utilities today is forecast-based. Utilities identify a need for new resources by forecasting future demand for electricity, and choose new resources, whether supply-side or demand-side, based upon a comparison of the forecasted costs of the various types. Despite efforts to improve the capability of forecast-based planning to identify and address uncertainty, uncertainty remains one of the greatest weaknesses of forecast-based planning. An increasing probability of discontinuous change in the coming decades only accentuates this weakness. An alternative approach is goal-based planning. Goal-based planning starts with an identifiable, measurable end-state deemed beneficial under all or at least most potential future scenarios. Utilities would then use forecast-based planning to choose the timing and content of paths to move toward the identified end state. The ability to observe progress toward the goal under the path chosen and adapt as necessary to continue progress toward the goal would inject an important learning loop into the process. This approach would also allow the nation's many utilities to pursue different strategies to reach their articulated goals, enabling learning from each other. Goal-based planning can facilitate alignment of stakeholders on a long-term vision and channel disagreement into dialogue over the best way to reach the desired end state with actual results to help resolve differing views.

The roots of this paper lie in a personal story. Somewhat like Richard Fulgrum, who learned everything he needed to know in kindergarten, I learned everything I understand about the power of goals from a computer game I used to play with my son. The game was called "Carmen San Diego's Great Chase Through Time," and involved joining various historical characters, such as Leif Erikson or Lewis and Clark, and helping them perform some task. The most memorable for me was the task of helping Christopher Columbus sail across the Atlantic. In the game, your destination was clear: an "X" marked the spot on the opposite shore. The program required you to set the sails and the rudder and then push "enter." At that point, as they say, "life happened." Or, to be more precise, wind and currents happened. Your ship moved but not always in the direction you had intended. After moving, you knew certain things. You knew, for a given spot in the ocean, what the winds and the currents were relative to your destination. You still knew where your destination was. You set the sails and rudder and tried again.

After numerous rounds of the game with my son, I came to understand what was so important about that destination on the far shore marked with an "X." Both my son and I knew where it was and could use the same quantitative and qualitative tools both to determine where we were relative to it and whether, ultimately, we had arrived. Further, because we could "see" how we had moved relative to the destination after a given set of sails and rudder choices, we understood our environment – the relationship between the actual winds and currents and our sail and rudder settings – better. The destination, and the quantitative and qualitative tools available to us to understand where it was and where we were in relation to it, were critical. Without those, we would have, literally, wandered the ocean.

After many years in and around utility resource plans, I realized that – unlike the game described above – these plans had no goal. The plans compiled and compared many varied forecasts to support a decision to do something, rather than to go somewhere. With looming carbon regulation and other changes that will stress our ability to forecast the future as perhaps never before, the time seemed right to evaluate whether goal-based planning might serve us better than our traditional forecast-based planning.

Although subject to many variations, current utility resource planning relies heavily on forecasting. Planning starts with a forecast of demand for electricity in the future. In simple terms, we identify existing relationships between such drivers as economic conditions and housing choices and electricity consumption. We then forecast future loads by assuming such relationships persist and inputting to the model assumptions for each of the drivers over the forecast period. We also forecast the state of existing resources over the same time period as the load forecast and determine whether a gap exists. We forecast the future costs of numerous resource choices, supply-side, transmission, and demand-side. These forecasts encompass everything from the price of commodities such as natural gas and coal traded in a global marketplace and the "costs" of environmental effects of various electricity conversion technologies that are not yet expressed in a control technology cost or other internal cost form. At this point, planners may well translate all of the forecasted costs to a net present value by applying a discount rate. Through this calculation, costs deemed to occur farther in the future appear much smaller than forecasted costs incurred in the near term.

We know these forecasts are uncertain, including the chosen discount rate. So, we commonly do a range of forecasts; perhaps a high, medium, and low. Or, we may assemble scenarios in which two or more of the ranges of uncertain forecasted costs interact. If the planning work is very sophisticated, it may allow certain variables to change in a Monte Carlo simulation. Nonetheless, at the end of the process, resources have to be chosen to fill the base case gap identified, and the predominant means of choosing them will be a comparison of the costs of certain choices against other choices. Once the utility has proposed and the Commission accepted a set of resource actions and the resources chosen are in place, the set of forecasts underlying the plan will typically be shelved and gather dust. Events will unfold; the utility, perhaps in consultation with stakeholders, will react to the events, and so it will go. The next plan will start over with a new set of forecasts.

Aside from a very human inability to predict the future, several specific tendencies hamper our ability to improve how we assess uncertainty in resource planning. They include:

Linear orientation. Forecasts of load or a given cost input to power production tend to be linear, with uncertainty expressed as different slopes. For example, most utilities produce a base load forecast and a high and low case that will match the direction of the base but with greater and lower angles, respectively. This linear orientation typically excludes both step changes and the possibility of the line moving in the opposite direction.

Sight restrictions. We must recognize an uncertainty to consider it. Unfortunately, what we perceive is usually limited by what we believe is relevant and that can exclude matters one or two steps further up or down the supply chain.

Implicit goals. Participants in the planning process may well have preconceptions about an acceptable outcome, based on beliefs regarding the desirability or undesirability of various resource choices and/or the contribution a given choice will make to the ability to reach some unexpressed end state. These implicit goals can exacerbate the sight restrictions as we are generally unwilling to go look for uncertainty about something we have decided we want or don't want.

So, uncertainty exists, we make decisions anyway, and then life happens. In my experience, we don't often look back at our forecast-based plans. Certainly, I have found no easy place to go for this type of retrospective analysis. And from a regulatory perspective, why should we? The decision whether a given resource action was prudent must be based on what was known or knowable at the time of the decision. What actually happens is irrelevant, or at least is supposed to be.

Almost everyone who has ever participated in a planning process resulting in a resource decision implemented at least ten years or so ago, can think of an example of how things did not turn out as forecasted in the planning process. The most notable example in my personal experience relates to the forecast-based resource planning process by which the utility I was working for at the time – Portland General Electric Company (PGE) – decided to close the Trojan Nuclear Generating Plant because the forecasted costs showed that it was less expensive to replace it than continue to operate it.

As I look back at this 1992/93 decision, I see two kinds of forecast-based issues. The first set falls in the category of things that we know are likely to be different actually being different than we forecasted, such as loads, fuel prices, and – in this case – plant operating conditions. By 2008, about 17 years after the forecasts in question, loads had grown at a rate just under 1.1 percent annually, just under the bottom of the low load forecast in the plan of 1.1 percent. Natural gas prices as of March 2009 were not far off the 2010 amount in the plan (real dollars): $6.09/MMBtu in the plan and $6.86 March city gate price from the Energy Information Administration (EIA). More interesting is the pattern of gas prices during the intervening 17 years, which ranged from a low of $2.74/MMBtu to a high of $9.49/MMBtu (nominal dollars). And, although it is impossible to know what Trojan's operating costs would have been, the experience of the late 1990s and early 2000s suggests that the plan's forecasts of Trojan's future operations and maintenance costs were perhaps too high, and that the forecasted cost per MWhr should have included a scenario under which Trojan adopted a 24-month refueling and maintenance cycle, as many plants did in the late 1990s.

The second set of forecast issues falls more in the category of difficulty in evaluating or even perceiving an uncertainty. For example, how the forecast-based plan addressed carbon regulation and nuclear externalities is interesting. The plan explicitly mentioned both uncertainties but, in a short paragraph, concluded that because neither would change the result, the plan would not attempt to value them. And perhaps most interesting in terms of an uncertainty not even identified, the plan does not mention the possibility of independent power production and competitive wholesale and retail electricity markets. In 1992, with the ink on the National Energy Policy Act creating exempt wholesale generators and authorizing the Federal Energy Regulatory Commission to order transmission service, the possibility of these markets and their effects on the West Coast were highly speculative.

Thus, the plan was done; the Oregon Public Utility Commission acknowledged it; and PGE closed the Trojan Nuclear Plant. As noted above, then life happened. In the first years, the benefits of closure were far greater than anticipated, as the opening wholesale markets and West Coast power surpluses combined to produce unbelievably low prices for power purchased to replace the plant's output. After building one natural gas combined cycle combustion turbine plant to partly replace Trojan, PGE postponed any further construction and relied on now-cheap market purchases. Then the West Coast power market crisis hit and wholesale market prices and natural gas prices reached levels unconceivable to anyone. The benefits of closing the plant became detriments, particularly if one assumed Trojan might have achieved the performance of many of the nation's nuclear plants, which became among the lowest cost resources. As the wholesale market began to settle down, concerns about climate change began to rise and costs associated with carbon regulation have become more certain. Long-term costs associated with nuclear power production externalities remain murky, at best, however.

So, was Trojan's closure the right choice? Or is that even the right question? How could we actually know whether it was the "right" choice or not, given that there was no goal at the time of the decision other than perhaps an implicit objective of "low rates" in the future?

The weakness with forecast-based planning is that it is entirely possible to make the "right" decision or even "least cost" decision, plan after plan, and still end up somewhere that one completely regrets. This is particularly the case in times of great potential change. Looking at the 20th century, Vaclav Smil recently called our record of energy forecasting, whether regarding broad consumption trends or the pace of commercialization and adoption of new technologies, an abject failure.[i] It is easy to imagine that the record for the 21st century might be equally, if not more, dismal because of the possibility for discontinuous change or, in other words, changes that fall off the trend line.

Changes that lay waste to our careful trend lines could occur in at least three broad areas: technology, policy and behavior/belief. We likely feel most comfortable with our observations of potential technological changes, particularly those that relate to how we convert various forms of energy into electricity and how we might store electricity once generated, knowing the "what" if not the "when." Nonetheless, surprises could easily occur here. Coming change in information technology is harder to observe but is nonetheless real and at a pace that could definitely surprise electric utilities, just as telecommunications utilities were often surprised by the information-based technology changes that swept through the 1990s. We now expect policy change around carbon dioxide and climate concerns but the coming decades hold the possibility of even more dramatic policy change as the world copes with a still-rising population, many of whom have extremely limited access to electricity, and economic shifts in what we make, where we make it and what we do with the stuff that's been made but is no longer useful. And interacting with both technology and policy is behavior and belief. Planners tend to assume that people will continue to act as they have before. Yet, in *The Tipping Point*,[ii] Malcolm Gladwell takes us through numerous examples of beliefs and behaviors changing like wildfire after reaching a certain point, usually not discernable before the fact.

When faced with change, we have two fundamental choices. We can try to fight the change, actively or passively by fervent hoping that it doesn't happen or goes away. We often call this "managing" the change. Or we can set ourselves up to embrace the changes. Goal-based planning enables the latter choice. In key respects, it is a time-honored choice. A good example close to home might be the electrification of Chicago at the turn of the 20th century. When Insull began, next to none of the city's structures were wired for electricity, people didn't use electricity in their lives or livelihoods – they had other ways of accomplishing everything that electricity could do, no coherent policy existed for regulating the provision of the service and technology was rudimentary at best. Within 20-30 years, the city was – for all intents and purposes – electrified. This did not happen because Insull forecasted that it would. It happened because he set out to lead.

The power of goal-based planning lies in articulating the goal. Before defining a goal, however, it is necessary to back up a step and talk about purpose. Purpose is, at a very fundamental level, <u>what</u> you are trying to do, separate and apart from the present <u>means</u> by which you are trying to do it. Purpose provides the boundaries for a goal. If one's purpose is too narrow – for example, tied tightly to the present means of doing what it is the organization does – then a goal formed within it will also be too narrow and unlikely to survive profound technological, political, and cultural change. Thus, goal-based planning for utilities requires us first to articulate a purpose for utilities that is broad enough to serve our objective of leading planning through decades of uncertainty and change. Although it is most meaningful for a group of stakeholders to share a dialogue on this and reach their own articulation, I will short-circuit that for purposes of this paper and suggest that one could articulate purpose for electric utilities as helping people apply energy in their lives and livelihoods. This purpose nicely encompasses activity in which utilities are already engaged: helping people with improvements to their structures such as insulation that lessen the need for electricity to condition the space comfortably. This is not providing electricity, nor even providing "negawatts." It is helping people apply energy, in this case by hindering energy flows, to achieve a desired outcome: comfortable space in which to live and work.

The goal, then, describes a picture of what it will look like if the utility is successful in helping people apply energy over the coming decades; how meeting that fundamental human need will be better than it is now. The center of the goal is the people and the goal must be from their perspective; i.e., what it will look like for them. Moreover, the picture should have sufficient detail that two or more people can hold a similar understanding of it. And it must be capable of expression through any number of indicators, ways in which everyone involved can know:

- Where we are now
- Where we are going
- How far we just moved, with the strategy most recently adopted

This definition rules out goals around rates, or cost, for electricity. Rates and the costs they derive from are inherently from the utility perspective, not a customer or people perspective. It is also highly unlikely that two people could hold a similar understanding of rates with some modifier such as "low" or "reasonable." And, for that same reason, it is impossible to determine where you are now, where you are going, and how far you just moved with a given strategy, toward a goal expressed in terms of rates. To return briefly to our earlier example of the planning process by which PGE closed Trojan, if the goal had been low rates in the future, some might have argued the Plan initially achieved that and to a greater extent than anticipated; then absolutely did not achieve that, as the West Coast power crisis hit; and then may or may not have achieved that, depending on how low one thinks "low" is.

As with purpose, the goal for goal-based planning is best found through deep dialogue among the stakeholders affected by it, drawing on their diverse perspectives of what certainties and uncertainties the future holds.[iii] We may not know what specific natural resource inputs to applying energy will cost decades into the future, but we can be certain that, with the current and projected population, competition for resources will remain a force. Therefore, better energy application solutions will likely be those that require fewer natural resources. Similarly, we may not know specific environmental policies that will emerge, but we can be certain that, as our knowledge grows of the effects of various manufacturing waste flows on our health and the health of other living things, policies will seek to limit those wastes. Therefore, better energy application solutions will likely be those that require less of manufactured things, including the conversion of primary energy into electricity.

A goal should transcend our current calculations of cost and conceptions of cost-effectiveness. Cost becomes relevant at the stages of "how fast" and "by what path." But history has so many times taught us that what we really value in the future are the decisions made for enduring community values. How many cities, in recent years, have expended billions of dollars to move major roads and freeways from their waterfronts, belatedly recognizing what providing the community access to the river or bay was worth? Each time, the community undoubtedly wishes someone had thought of it sooner, had seen the value and chosen a path, at a cost, that would have preserved it. How a community converts and delivers energy to its purposes implicates similar enduring values. Goal-based planning supports us asking ourselves: what would make it safe to assume that, fifty years from now, people would be happy we made this decision?

Again I will take a short cut here, for purposes of completing the story of why goal-based planning is better suited to the 21st century than our familiar forecast-based planning, and suggest a goal that utilities could use. It is the following:

People in the communities we serve will apply energy to meet all of their individual and business needs at the same pace that nature provides it (no faster than the rate of regeneration of what is renewable), with no more waste products than the rate at which the earth can recycle, absorb, or render them harmless, and in a way that strengthens their connections with this place and with each other.

With the goal in hand, the next step in goal-based planning is to assess where you are. This requires identifying at least the initial set of matters relevant to create a shared understanding of elements in the picture. Working with the articulation above, that list of matters might include for a given community:

- The amount of energy available for application, whether directly or through conversion
- The nature and amount of waste products produced in the community's application of energy
- The amount of energy required for various individual and business needs
- Its sense of connection with the sources of the energy it is applying
- Its perception of shared responsibility and reward for the energy it applies

The planning process then consists of assessing the current and ideal states of the relevant matters and identifying and evaluating paths to narrow the gap between the current and ideal states. It is here that forecast-based planning re-enters the picture, along with the familiar decision-making tools of effectiveness and cost-efficiency, for the planning process participants must decide:

- How much of the gap to address
- Over what period
- Through what means

While it remains probable that various participants in the process will see these matters much differently and disagree, goal-based planning offers some discipline on these disagreements that forecast-based planning does not. That discipline is the knowledge that it will be possible to check at least the effectiveness of the strategies chosen in decreasing the gap between current reality and the goal. It will also be possible to compare utilities to each other in terms of the effectiveness and cost-efficiency they have chosen to work toward their goals. With forecast-based planning, disagreements are essentially incapable of resolution because there is no goal against which to compare effectiveness or cost-efficiency.

Goal-based planning does not eliminate uncertainty. What actually happens will be different from what we forecast will happen. The value of goal-based planning is that it lessens the chance we will be seriously unhappy with what we have done. Utilities planning according to a well-formulated and articulated goal will, albeit with many wanderings, move closer to that goal over time. They will learn what works and by their intention to move closer to the envisioned future, increase its likelihood of occurring. If the goal of planning at the time of PGE's decision to close the Trojan nuclear plant had been to develop a resource portfolio with which its customers and communities had greater comfort, then the closure of Trojan likely moved PGE closer to the goal.

Forecast-based planning is a bit like wandering the ocean, in hopes of landing in a good place. As the pace and scope of change increases, however, utilities are increasingly unlikely to do so. Moreover, by wandering about, utilities and their stakeholders will forgo the opportunity to influence, or even to lead. Goal-based planning gives us a destination and a shot at leading. In a carbon-constrained world, I know what I'd rather be doing.

---

[i] Smil, Vaclav. *Energy at the Crossroads*. MIT Press, Cambridge, Massachusetts, 2003.

[ii] Gladwell, M. *The Tipping Point, How Little Things Can Make a Big Difference*. Little, Brown and Company. Boston, 2000.

[iii] California recently engaged in a process with many of the elements discussed here: its California Long-Term Energy Efficiency Strategic Plan. Through a collaborative process, the California Public Utilities Commission adopted four specific goals to drive future energy efficiency planning. It is not clear yet, however, how the Commission will implement this strategic plan in its decisions regarding individual utility proposed energy efficiency programs and plans.