Analytic Intelligence: The Core of the Utility Smart Enterprise

ORACLE WHITE PAPER | SEPTEMBER 2015

Introduction

A “smart” enterprise optimizes the entire operation of a utility—from the power plant to the customer to meter to cash. Using available resources, the smart enterprise extracts insights from data provided by multiple sources to drive efficiency, grow customer satisfaction and bolster safety records. This is the promise behind enterprise analytics and the hoped-for end goal in an analytics initiative. The right analytic solution should not only showcase early and quantifiable results but also enable budgetary flexibility through business impact in both cost and revenue drivers. Analytics, however, have often been viewed as peripheral to the core business of utilities. The disconnect is real: dashboards feel ancillary to the core business of delivering megawatts and keeping the lights on.

This paper explores why this perception is unfounded and shares a new paradigm for the industry to consider when seeking the goal of an optimized enterprise. Establishing a common and structured framework for determining business value is critical in creating the conditions for a successful analytics initiative. Without this context, analytics becomes a bogeyman of failure—ill-defined problems, KPIs that measure the wrong values, insights without action—and the utility continues to work in the dark without the insight provided by data. In these cases, the problem wasn’t analytics; instead, the analytic initiative was not set up to succeed in an environment of quantifiable successes and calculated business benefit.

How, then, can we avoid the missed expectations of analytics strategies attempted in the past? The first step is to explore a thorough discussion which re-contextualizes advanced analytics vis-à-vis traditional business intelligence and traditional consulting services to identify the flaws created by these approaches. Next, we review what defines a modern and successful analytic deployment through three fundamental lenses which define our new paradigm: process context, time responsiveness, and organizational culture.

Advanced Analytics Moves beyond Business Intelligence and Consulting

Advanced analytics drive operational efficiency and innovation through the integration of business functions and the streamlining of existing business processes. Business intelligence is in many ways a precursor to advanced analytics. While the reporting functionality of dashboards—traditional business intelligence—illuminates the inner workings of an organization to some extent, dashboards are static in nature while operations are inherently ever-changing. A dashboard of performance metrics for billing operations only captures a snapshot in time of what the business values are. Additionally, it requires the viewer to translate the report into action with no specifics on what, how, and when to act.

By contrast, while dedicated operational consulting projects do generate specific, action-oriented recommendations, the entire process—from the collection of data to the analysis and the refinement of the outputs—is manual and time-consuming. The smart enterprise consists of many repeated actions, and the ability to optimize operations pervasively must be done though the automation of the consulting process. Advanced analytics do exactly that—take in data, prepare the data for analysis, analyze the data, generate insights that drive action and then integrate feedback to the algorithms—all in a repeatable, ongoing fashion.

Here we examine how advanced analytics enhances the smart enterprise operations through three lenses:

» business process,
» time, and
» the organization.

Through the business process lens, we consider how different analytic techniques impact operational efficiency and propose a model for how to systematically enhance operational efficiency. Through the time lens, we look at specific scenarios where analytics enhances response time, supporting enterprise-wide strategic commitments to utility customers as well as regulatory compliance. Finally, we propose a qualitative perspective on the nuanced differences across the organization, and the need for targeted insights by function, through an organization lens framework.

Business Process Lens: Manual vs. Automated Analytics

One of the fundamental value drivers for analytic adoption is the ability to increase efficiency. Organizational efficiency is the foundation of bottom-line management as part of responsible corporate financial planning. As supported by many operations research studies, achieving this workforce efficiency is done through standardization. But what does standardization mean in the context of analytics integration? In what ways do analytics drive efficiency?

In establishing a framework to understand the impact of analytics on business processes, the two relevant axes are analytic customization and the complexity of the analytic operation. One measure of analytic customization is in the amount of resources or man-hours that go into designing and performing analytics. Similarly, the quantifiable proxy for operational complexity is the sophistication and volume of the data involved. Analytic customization is theoretically a continuous space, but the real-world established practices of analytic customization can be broadly grouped into three buckets:

» dashboards and reports,
» ad hoc search, and
» operationalized queries.

The number of man-hours required to generate each analytic product increases from dashboards to customized consulting. Similarly, business intelligence (BI) dashboards often capture people and performance; the data involved in monitoring billing exceptions addressed is actually far less than the data involved generating the billing exceptions in the first place. Generating exceptions accurately and comprehensively requires significantly more data and time to analyze.

The natural state of an organization (or the corresponding amount of resources that go into accomplishing various levels of operation complexity) can be considered a type of "efficiency frontier." Machine learning involves a group of advanced analytic methodologies that enable more operationally complex activities to be completed with fewer resources. Machine learning enables this by not only managing the volume and sophistication of the data more efficiently but also intelligently automating the analytic actions that would otherwise be done manually.

To put it another way, there is a theoretical boundary for how much customized analytics is required to produce useful results for more sophisticated business operations or processes. For example, a BI dashboard is not going to be sufficient for forecasting: a customer analysis is necessary for that initiative. The more complex the operation is, the more customization is necessary for useful analytics.

The graph below illustrates the tradeoffs between operational complexity and the amount of analytic customization necessary. If we assume that operational standardization is constant, then for a given level of operational sophistication, the efficiency frontier is where we will find the analytic solution that requires the least amount of customization necessary to satisfy the needs of that business process. The natural position of where real scenarios fall within this framework is one where simpler business processes require less customized analytics. Note that the actual slope or exact shape of the line is not relevant in itself. It simply illustrates the impact that more sophisticated statistical/analytic methodologies have, and how they influence the line. For example, machine learning is a more sophisticated statistical/analytics methodology, and with it, you can get more useful results with less customized work for more complicated operations. The "analytical intelligence" of machine learning facilitates a lot of the manual custom analysis, and increases the slope of the efficiency frontier line.

While the gauge of an organization's efficiency frontier is a relative one, the enterprise always has a natural benchmark within itself through time. With the advancements in technology, an organization should see a shift in its efficiency frontier to the left (i.e., for a given business process, there should be less analytic customization).

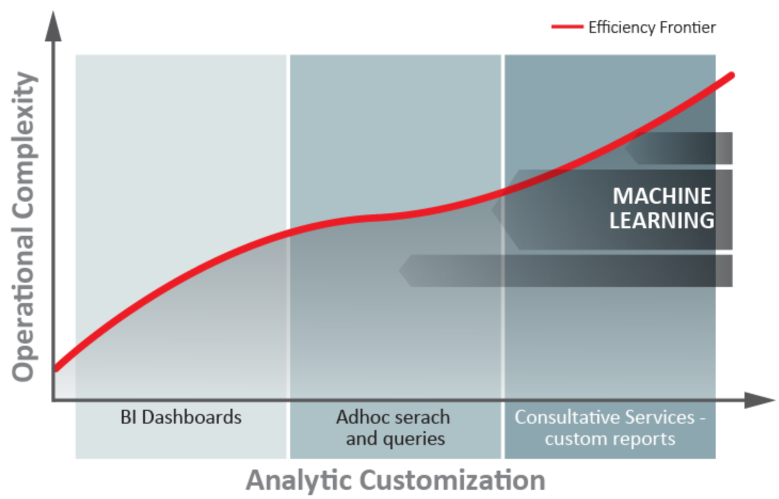

Figure 1: This efficiency frontier graph illustrates an organization's operational boundaries with respect to analytical complexity.

Part of this technical advancement is in the form of advanced analytics. Machine learning is an advanced analytic technique that has many applications within the utility to move its efficiency frontier. While the concept of machine learning has been around for a long time in the context of analytics, the specific, effective application of machine learning to utility use cases is relatively new and not yet pervasive. Machine learning or other advanced statistical methodologies, in combination with deep domain expertise, can successfully shift the utility's efficiency frontier.

One excellent use case for machine learning is in load profiling applications. Here is an example of an Oracle DataRaker deployment currently in use:

Historically, load profiling has been a semi-manual exercise performed on an irregular schedule dependent upon regulatory needs and customer demographic shifts. Despite the infrequency of this analytic exercise, it is the basis for many important processes for the utility and therefore should be performed using the most accurate, results-producing methodologies. One primary use case for a load profiling output is to optimize the rate structure. In this specific case, the output of the load profiling exercise is used to estimate monthly meters at higher resolution for forecasting applications. The beneficial implications of higher-resolution information based on monthly data are many, ranging from more accurate estimates for monthly metered customers to extending the life of existing monthly meter assets and getting more use out of higher resolution assets.

Load profiling is a fairly sophisticated business process with ever-changing steps depending on the data. In short, it involves understanding not only what groups each meter belongs to but the number of groups in total. Then the actual load profile (which is a continuously changing metric) must be generated for each group. If done manually, significant resources can be spent understanding each customer or premises attributes and its relevancy to predicting the load profile group that the monthly meters belong to and hand picking the right ones. In this Oracle DataRaker application, K-means clustering was used to select the number of groups that the entire population falls into and assign each meter to the appropriate group. This in and of itself reduced the most significant portion of the manual work done in load profiling. Additionally, it enabled the processes to be done in an automated fashion, which then allowed for more requirement evaluation and corrections made to the load profiles and groupings.

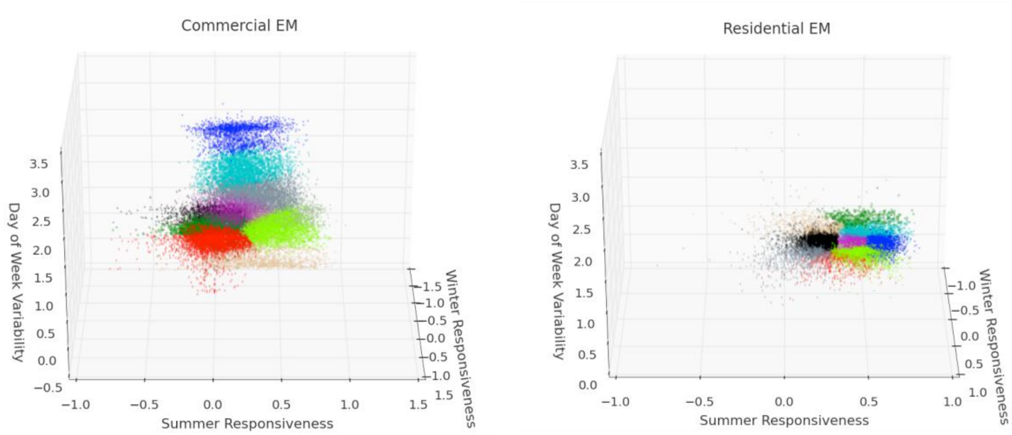

Figure 2: Example of machine learning results and clustering results for electric meters. Each point represents the average Summer Responsiveness, Winter Responsiveness, and Day of Week values for all meters with a specific combination of demographic variables. Clustering results are denoted by color, with each color representing a different customer segment.

Leveraging machine learning initiatives can shift a utility's efficiency frontier by minimizing costly manual intervention and achieving greater business value through automated analytics. In our "Machine Learning for Smarter Load Segmentation" white paper, we showcase in more detail the complexities of applying advanced analytics to utility use cases and the effect of the application of domain expertise with subject matter expertise.

Time Lens: Quick Response Increases Financial Effectiveness

Another benefit of successful analytic adoption that has direct impact on top-line generation is the organization's timeliness in responding to issues. In other words, the critical aspect of analytics in an enterprise application is the timeliness of the analytic insights. Particularly in the realm of operational analytics, having a quick response time relative to the occurrence of the event of interest is an important component of the insight's relevancy and also financial impact to the organization.

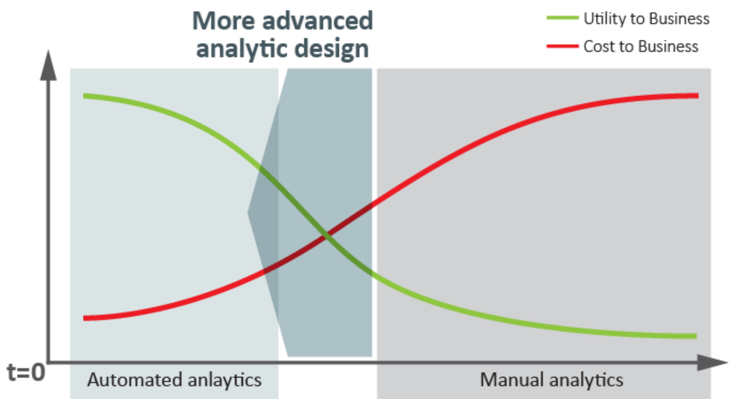

Figure 3: Usefulness/cost vs. time: Faster response time drives operations closer to the efficiency frontier.

A piece of insight's usefulness is inversely proportional to the consequence or cost to the business; more simply put, usefulness is high when impact is high. In abstraction, as time passes after the occurrence of an event, the usefulness of a response and reactive action decreases and the consequence or impact to the business increases. A frequent example of this within the utility is broken meters. If the broken meter is discovered immediately after it malfunctions, the revenue collection consequence will be minimized if that piece of information (i.e., the discovery of a broken meter) is communicated in a timely manner. In the example given later in this section, the information collected on orphan meters is only relevant until the next meter is installed; after a period of time, that particular piece of information becomes useless.

The concept applied to this example is often summarized in the "break to fix" Key Performance Indicators (KPI). Many things can impact how an organization operates with respect to relevant events for core activities. Operational analytics can influence the organization's responsiveness for key actions. Again, using the example of broken assets, manual ongoing review and analysis will result in significant delays to response time. Partial automation of analytics can shift an organization's operating zone to the left (as indicated in Fig. 3). However, to truly systematically increase the organization's efficiency and effectiveness, analytic insight needs to be provided as close to the event of interest as possible. In order to do this, the analytic solution needs to be iterative, and it needs to be implemented with closed-loop automation. Closed-loop, automated analysis enables seamless processing of new information as it is provided and provides the most timely insights in the most efficient way.

This concept is best shown in two different utility use cases. The first and most direct example of the concept is illustrated by an Oracle DataRaker analytic solution in use within another utility by its meter services group. A series of meter operations applications were used to enable more efficient operations of the utility's meter services, and this had an immediate impact on the utility's break-to-fix times. Three snapshots of the group's break-to-fix performance were taken: before the deployment of Oracle DataRaker, shortly after and a year later. The shift to reduced break-to-fix times was immediately evident and, over time, the impact became more pervasive over the entire meter population.

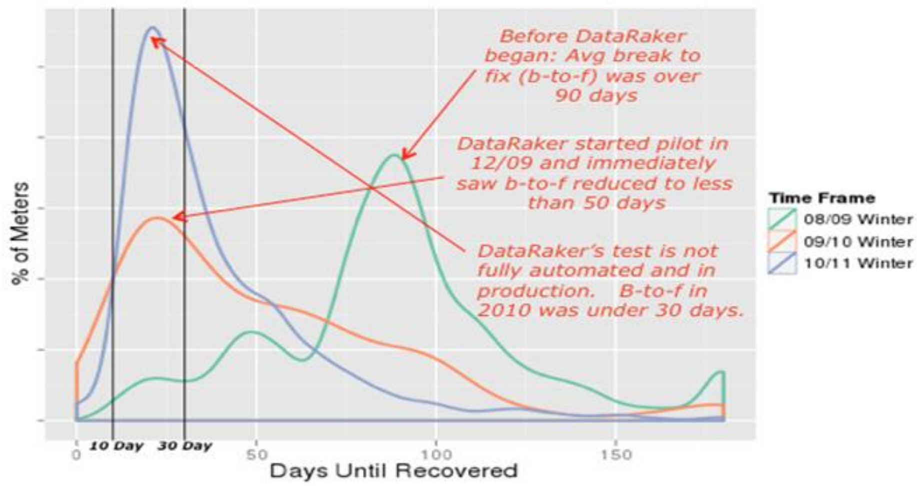

Figure 4: With Oracle DataRaker analytics, this utility saw a reduction in break-to-fix time of more than 50% for the bulk of its malfunctioning meters.

Dealing effectively with orphan meters is a second utility use case that highlights in more detail how analytic automation enables responsive timeliness. Orphan meters are defined as physical devices that are lost in the field. The situation is instigated by a meter-swapping event such as an advanced metering infrastructure rollout or the replacement of a broken device. Often, the new meter is misconfigured and never successfully communicates its meter reads to the head-end system, even though the premises continues to consume energy. The timeliness of identifying and recovering the meter is crucial to its financial impact on the organization. The longer it takes to identify the lost meter, the longer the consumed commodity remains unbilled. This particular issue is also analytically challenging to discover, and so many utilities without the right analytic solution often omit these meters altogether. This, in turn, incurs wasted asset costs. Also, because the consumption and event data are captured before the meter is swapped, if the new, non-communicating meter is swapped again, then that new meter is forever lost. This adds a secondary timeliness component to identifying these orphan meters.

For this use case, Oracle DataRaker successfully identified an analytic approach leveraging a Gradient Descent Classifier technique to most accurately and efficiently identify orphan meter scenarios. This analytic exercise is performed daily with automated processing that generates a list of potential orphan meter candidates. The orphaned meter is detected almost immediately, minimizing unbilled consumption cost in addition to the risk of forever losing the meter asset.

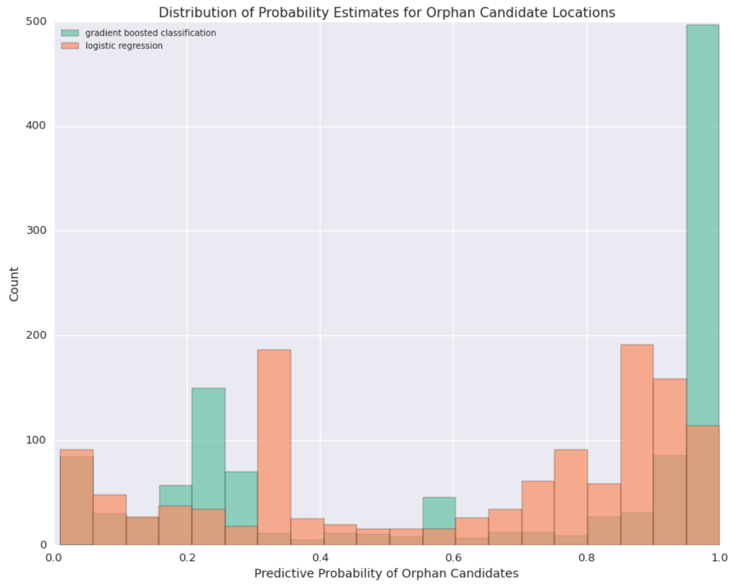

Figure 5: An example of results from orphan meter analysis. Analytic feedback was able to enhance output performance by better methodology selection.

The predictive probability is an internal assessment of a methodology's effectiveness which, in turn, impacts the time to detection. In this sample graph, we see the gradient boosted classification methodology is significantly more effective than the alternative logistic regression methodology. The enhanced performance can be the difference between identifying the orphaned meter before a subsequent meter swap happens and forever losing that meter asset.

While reducing break-to-fix times clearly shows quantitative results, there are also qualitative results in terms of increased customer satisfaction. In addition to the benefits defined above, finding orphaned meters sooner also reduces the incidence of large catch up bills for customers. This, in turn, also increases customer satisfaction. Using analytic automation to more quickly and effectively identify orphan meter scenarios not only reduces the cost to the utility from a quantitative perspective in terms of asset recovery, but also in a qualitative way through better customer service.

Organization Lens: Infusing Analytics into Your Culture

There are countless variations in analytic insights: the different types of business operations vary tremendously between different functional groups and analytics must reach the long tail of operational knowledge workers that work in these different business operations. It must focus on volume effects for optimized margins through functional and application reach. The types of analytic insights a knowledge worker needs compared to the needs of someone in management or an executive are different. Each operational knowledge worker needs a slightly different insight, as well, depending upon the operational process and business function he or she supports. While a handful of dashboards can satisfy a sales executive regarding revenues and various segmentations of revenues, both the insight and mechanism of how an analyst in meter services receives the insight is drastically different than how it is received by an operations worker in network management. The frequency that a dashboard can be used within the enterprise is high while the specific insight related to broken meters is relevant to only few.

However, there are many instances of similar-level insights needed throughout the organization, given that the bulk of head count within a successful organization is not in management. Yet, to maximize the impact of analytics throughout the organization, the right insights must reach the right people. This requires a combination of the right technology offering with some organizational change management to ensure adoption. Only then can the enterprise realize the synergistic effects of more efficient and accurate performance.

Analytic pervasiveness also applies horizontally across the different business functions within a utility. Analytics does not only apply to the meter-to-bill work flow but also in distribution operations and customer services. In our "Mastering Smart Meter Advanced Analytics for Water" white paper, we examine the full suite of analytic opportunities across a water utility leveraging all sources of sensored time-series data.

Adopting advanced analytics both vertically—from analysts and managers to operations, network and customer service employees—and horizontally across the utility's business units allows the entire enterprise to directly benefit, albeit in different, though synergistic, ways throughout the business. As this approach filters through the enterprise in both directions (vertically and horizontally), the overall benefits to the utility's efficiency and effectiveness will continue to grow.

The Opportunity Cost of Not Investing in Analytics Core Technology

Successful implementation of analytics into the organization's core has three attributes: operationally integrated analytics based on the successful use of advanced statistical methodologies; automated, advanced analytic processing; and breadth and depth in analytic reach. These three elements drive a solution that satisfies both in function and timeliness and effectively scales throughout the organization. This in turn manifests in several quantifiable benefits for the enterprise, including quick implementation and time to value and the ability to prove its own worth within a specific budgetary period (in the case of a regulated utility in North America, this would be the time period of a rate case). Tactically, this involves carving out Opex funds through operational efficiencies and recognizing more revenue.

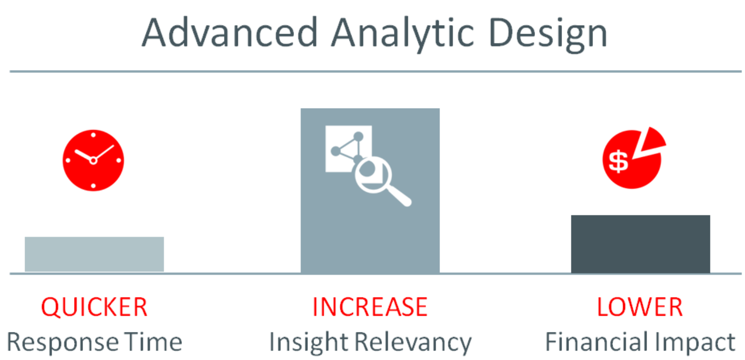

Figure 6: Establishing a comprehensive analytics strategy will enable all organizations across the enterprise to operate seamlessly and efficiently for net quantifiable results that impact both bottom-line management and top line revenue growth.

Another way to ask the question is, in the absence of a successful analytic solution, what is lost? The benefits of this analytic adoption can be evaluated using four objectives, with each satisfying a subset of the analytic application of use to the utility. These are outlined below. The first two objectives, revenue assurance and preventative maintenance, are hard dollars that contribute to the payback of the analytic solution. These are line items that free up more Opex spending within a predetermined budget. The second two, customer satisfaction and safety, are soft benefits that are theoretically priceless but, when quantified, further support the argument for analytics.

Objective	Description	Structure	Relevant Use Cases
» *Revenue Assurance*	» *Ensuring top line generation is maximized*	» *Reducing write offs from commodity lost through distribution inefficiency and theft* » *Decreasing collection times and loss from processing inaccuracies*	» *Billing operations* » *Line loss monitoring* » *Theft detection* » *Rates optimization*
» *Preventative Maintenance*	» *Ensuring bottom line maintenance is optimized*	» *Reducing Capex and associated Opex expenditures* » *Reducing Opex through more efficient operations*	» *Transformer failure detection* » *Meter failure detection*
» *Customer Satisfaction*	» *Synergistic effects from pervasive analytic adoption*	» *Enabling the utility to deliver better service as measured by CSAT and subsequently impacting the utility's cost of capital*	» *All use cases*

Revenue assurance is primarily driven by operational efficiencies (reduced or shifted headcount) and recovering what would have been write-offs stemming from billing estimates, theft and distribution loss. Operational efficiencies are achieved through automated manual analysis with an analytic platform and services. The ability to recover what would otherwise be lost revenue can be attributed to more accurate analysis performed by the analytic platform and services.

Preventative maintenance has two contributing sources. The premise of this assumption is that the analytic platform and services enable enhanced performance of existing assets and therefore extend the life expectancy of these assets. Financially, this translates to a longer depreciation schedule which, in turn, reduces the annual Opex associated with depreciation. Extended asset life also means a reduction in Capex spend and associated interest paid on that investment.

There are also synergistic effects of analytics adoption throughout the organization. As a result, all aspects of the utility's services delivery are perceived to have improved, thereby increasing customer satisfaction as well as the utility's Customer Satisfaction Score (CSAT). The utility's increased CSAT, in turn, lowers its risk profile and correspondingly increases its bond rating. A higher bond rating nets a lower cost of capital and reduced Opex for the utility.

Conclusion

While there are many ways to quantify the value of analytics, its true value to a utility enterprise becomes much clearer when viewed through the three lenses of business process, time, and the organization:

- Viewed through a business process lens, the use of more sophisticated statistical/analytics technology provides the utility enterprise with more useful results with less customized work necessary for the more complicated operations.
- Viewed through a time lens, an iterative analytic solution that is implemented with closed-loop automation enables seamless processing of new information as it is provided and yields the most timely insights to the utility in the most efficient way.
- Viewed through an organization lens, the combination of the right technology offering with some organizational change management allows the utility enterprise to more fully realize the synergistic effects of more efficient and accurate performance.

Utilities are entering a new era of customer expectation and technological innovation, and analytics is a core requirement for the truly modern digital utility. Those with pervasive analytic adoption will be able to internally comprehend and act upon inefficiencies and operational pain points to not only maximize existing assets but provide the best service possible to their customers. For retailers in unregulated markets, analytics can be the differentiating factor for top-line growth through enhanced service and higher margins. And for both regulated and unregulated markets, analytics also enables a level of transparency that will enable each utility to respond to changes and issues faster and delivery services with fewer incidents.

Contact Us

For more information about Oracle Utilities solutions, visit **oracle.com/industries/utilities** or call +1.800.ORACLE1 to speak to an Oracle representative.

ORACLE®

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oracle
facebook.com/oracle
twitter.com/oracle
oracle.com

Integrated Cloud Applications & Platform Services

Analytic Intelligence: The Core of the Utility Smart Enterprise
September 2015
Author: Oracle Utilities
Contributing Authors: Kate Rowland, Ilyssa Norda