Automating Intelligence

Developing an Organizational Culture to Maximize your Analytical Initiatives

Thomas Casey, Consulting Director, Analytical Solutions

06.18 EB10128

TERADATA.

Table of Contents

List of Figures

The Real Complexity Behind Analytics

The analytical world we live in today seems very complex, and in many ways, it is. The "Automating Intelligence" framework attempts to minimize this complexity by providing a structured approach for delivering analytically-enabled interactions within your organization and presenting guiding principles on how to advance such an endeavor. Limiting the complexity and putting into context these key components, however, is not the only factor in ensuring success. Operationalizing analytics requires a fresh perspective, an innovative spirit, the ability to challenge current thinking, and more importantly—a willingness to change. It is somewhat ironic that so many initiatives, aimed at leveraging keen insight and advanced analytics, fail so often because people are unwilling to follow where the analytics take them. Part of this reason is that the enablement of analytics is not a static process. There is no "magic statistical formula," as driving analytical solutions must remain iterative to enable a process of continuous improvement. According to James Guszcza, "by engaging in a snark hunt for the perfect model and by striving for impractical degrees of accuracy, statisticians sometimes sacrifice the benefits that could result from models that are imperfect but still useful."[1]

It is impractical to think there is a perfect answer to so many questions that analytics are meant to solve. Putting value on doing, learning, and improving requires a mindset where failure will happen. Nevertheless, in that failure one can learn and change the course of action. Leveraging an exploratory approach to analytics, however, is not an open-ended endeavor. As shown in this framework, we start with the data mining process which is based off key hypotheses that narrow the focus of the effort versus a "just look to see what I can find" mindset. This seems to be a dichotomy of sorts since we are asked to explore and "think outside the box" on one hand, but on the other, we need to be practical in that our outcome must ultimately lead to a viable business solution—as quickly as possible. Success in using this framework to drive operational efficiency requires developing a culture that encourages innovation but ultimately needs to lead to a reliable, working solution. Those objectives are somewhat at odds with one another in most organizations and is a major hurdle to achieving the goals put forth by this approach. To develop this culture; however, you must first look at it from two different organizational perspectives.

Preparing Your Organization to Embrace an Analytical Culture

At most companies, the IT organization and business end users that they support are at odds with one another. This is not because they do not like each other personally, nor are they collectively trying to do a poor job. Rather, they are at odds with one another because they are incented (and thus, motivated) to strive for diametrically opposed success criteria.

1 http://deloitte.wsj.com/cio/2012/06/06/5-reasons-analytics-projects-can-fail/

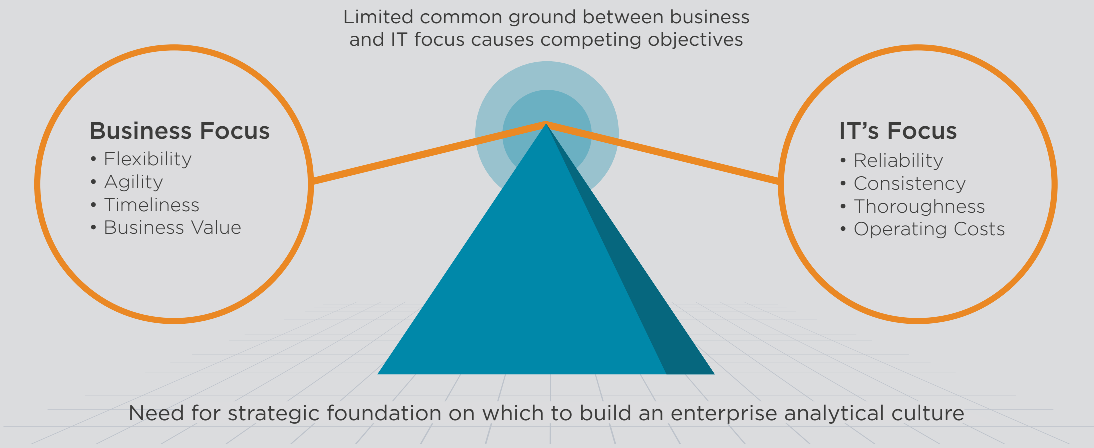

Figure 1—The Business vs. IT Conundrum

For business end users, competitors are always beating down the door and this constant pressure to deliver value makes "time to deliver" an ever-present companion. IT organizations, however, are much more conservative in what they do since they need to develop and manage solutions that are always working the way they were designed. Here is a visual of how the business' and IT's competing objectives are at odds with each other.

There is nothing inherently wrong with the focus of either group of constituents. However, there needs to be a cultural change to enable an environment of success. This is never easy but the following outlines one possible set of steps in driving towards this aim.

Step 1— Define/Refine Your Own Version of the Automated Intelligence Framework…and Make It Actionable

Our series of white papers provide a significant level of detail around the three strategies of the strategic Automated Intelligence framework. These provide an outline with several concepts to consider, as well as guiding principles for each strategy. Although this framework provides a practical set of analytical strategies, every deployment environment is different and there are several unique aspects to how this framework can be applied for a specific organization's needs. Different business priorities, adopted technologies, internal skill sets, budgeting constraints, etc., are all things that should be considered in the context of an organization's strategy.

Nevertheless, the framework in some form must exist to provide a foundation from which to operationalize analytics. Thus, it is imperative that consideration of all aspects of what has been previously outlined is incorporated when developing a thorough strategy. This includes the following:

Industrialized Analytics

Identify all the operational processes where analytics are or could be operationalized and assess your current methods and the effort required to move towards where you want to be. Be sure to take into consideration the following:

- Tie data mining activities to a specific set of outcomes and set boundaries on what they are looking to solve. There is an axiom, "you give an analyst six hours or six months they will take it." The meaning is that there is always a new pathway to explore with analytics, so one is never "done". Ensure that every data mining activity is designed to explore a specific purpose, define the parameters of that business purpose, and most importantly, give analysts a time limit to come back with a suitable result.
- View Artificial Intelligence (AI), not as a goal, but as a journey. Start with leveraging user-maintained models and expert rules—and design these to be data-driven processes (vs. having the inputs hard-coded). There are many places in any process where analytics apply.[2] Yet, machine and deep learning techniques require data, time, and training before models can be autonomous. A great deal of benefits can be achieved through managed models and not all need to reach the point of self-learning. Learn and evolve in your analytical journey.

2 See Automating Intelligence: Industrializing Analytics at Enterprise Scale

- Hire and train your analysts and data scientists to be creative, business inquisitive, and mathematically inclined (or balance the team with people of different skill sets). There are a disproportionate number of people in the analytical fields that focus on getting the "right" answer and far too few that have the creative and pragmatic instincts in determining how to tackle a business problem effectively. Analytics is a learning process so you need to build an analytical organization that knows that you need to do something to learn from it. Do not let "perfect" be the enemy of good.

Multi-Dimensional Personalization

If personalization is not a prevalent strategy within the organization already, pick a use case where personalization will make a demonstrable improvement in how you interact with a group of constituents. Typically, an organization will focus on some type of sales, marketing, or customer care interactions as increased value attributed to customer engagements tend to be easier to quantify. When developing a personalization strategy for this engagement process, ensure to adhere to the following:

- Confirmation that any personalization solution is designed to be channel agnostic from the beginning—regardless of the initial use case. Your interactions need to present a consistent "personalized" experience and incorporating different methods based on channel violate this premise. Thus, the scale of any personalization solution should begin with an enterprise view.
- Provide the flexibility to personalize across multiple dimensions. Personalizing on individual customer traits is relatively straightforward (in concept, if not in execution). However, one should and can leverage various segmentation strategies to personalize based upon groupings that, although not necessarily on an individual level, are able to support an intimate interaction in a way that ties to other business strategies (maximizing profitability, expanding customer relationships, etc.). Through this, you enable the eventual automation of your personalization using analytics and/or business-defined rules.
- Design a personalization method that enables you to personalize anything within the interaction—even if you cannot or do not want to personalize these elements today. Once you set arbitrary limits on what you want to personalize, you have limited your ability to successfully craft a one-on-one interaction. The mechanism to execute this personalization can happen across any dimension and can even vary from interaction to interaction in how these values are changed (see previous point). Nevertheless, do not limit what can be personalized or you limit your ability to support future requirements that will undoubtedly come. Whether these values are pre-defined or used as input to a real-time decision on how to react, only by starting with a design that can address any aspect of an interaction are you able to build a solution that meets the scalability and extensibility required to make personalization a competitive advantage.

Data Management Solutions for Analytics

Nearly all organizations of any significant size have undoubtedly invested in one or more data warehouse environments. Almost none of these categorize them along the lines as outlined in the Automated Intelligence framework to properly evaluate and position them with the correct context for enabling analytical goals.[3] Using the strategies outlined in this framework, each environment should be categorized based on analytical use cases as well as the following sub-strategy areas:

- **Data Sourcing**—What type(s) of data will be sourced onto this platform? What mechanisms are needed to source data into the environment based on latency needs? Is self-service loading enabled on this platform? How will data be transformed (if at all)?
- **Data Delivery**—How will data be delivered to a channel or another environment? What are the conditions where access will be direct, delivered (as a data feed), or some combination of both?
- **Data Stewardship**—How will data be stored? What security measures will be in place for this data on this platform? Is the schema necessary on-read or on-write?
- **Data Quality**—What measures will be taken to ensure data quality? What methods will be applied on this platform to fix data issues?
- **Meta Data**—How will data be categorized and tracked? How will users know which data elements are stored and where?
- **Master Data**—What is the strategy for integrating and enhancing all the elements required to define a robust customer master record? How will other reference data be adequately managed as part of this strategy?

3 See Automated Intelligence: Defining Data Management Solutions for Analytics

In addition to the above, there are three additional considerations with respect to the Data Management Solutions for Analytics (DMSA) strategy:

- Make a data "wish list" with business users. Determine the data that is available to drive analytics and personalization; but more importantly, identify the realistic data elements that are missing within the organization and develop a strategy to procure them. This does not necessarily mean that you will absolutely source the data, but you should have a working list of data elements the business says they need and the business benefits they will drive.
- Evaluate cloud-based deployments but do it from an infrastructure-only perspective. There are benefits to users in that these deployments can grow and shrink more rapidly than on-premise solutions (which may or may not include private clouds). Not only might this provide some flexibility as a test/development environment but also might allow users more flexibility in how they integrate with other third-party solutions (if it is a vendor-specific deployment). These come at a cost, however, and there is nothing inherently magic in using "cloud" as a strategy to get around doing something else that makes better corporate sense.
- Ensure that you have a sound Research and Development strategy for users to develop, hypothesize, test, and evaluate new ideas with analytics and data (see steps #2 and #3 for additional details on this point). This may include leveraging cloud-based deployments to support this strategy.

For each of the three strategies outlined (Industrialized Analytics, Multi-Dimensional Personalization, and DMSA), an architectural plan needs to be established and a conventional methodology for doing this is to break the architectural components into at least three groupings (there may be others but this is the minimum typically used):

- **Process Architecture**—the ideal state process of what is being optimized that identifies all the components from beginning to end, and the entities that will be incorporated in the process flow
- **Data Architecture**—the data necessary (include specific descriptive variables and measuring metrics) required to support, measure, and evaluate the end-to-end process
- **Technical Architecture**—identification of specific technologies and how they will be used in support of the various process and/or data architectural components

Viewing this visually, the framework can be represented by the following hierarchical graphic:

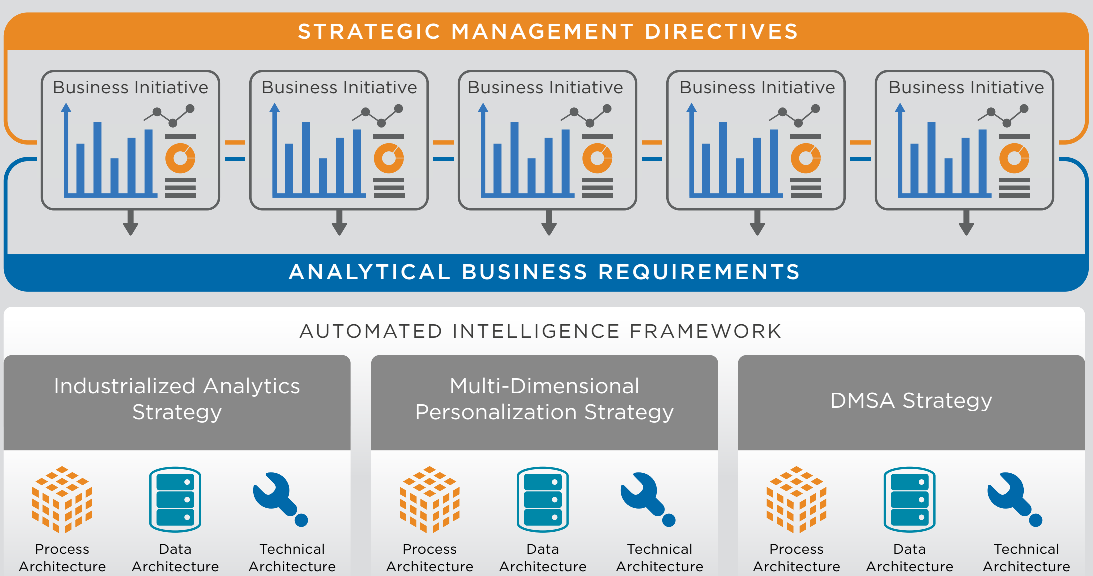

Figure 2—Strategic Operational Framework Hierarchy

Business initiatives with an analytical component, driven by strategic management objectives, produce a series of related business requirements. By establishing an analytical framework, one has a common set of strategies (i.e. guiding principles) on how to enable new business opportunities in their use of analytics. Executing on these business strategies, however, requires not only a framework but also a methodology for exploring opportunities in a timely and dynamic matter.

Step 2—Define a Robust Discovery Methodology...So the Business Can "Fail Fast"

A comprehensive framework sets the foundation, but one needs to have a flexible methodology that allows users to test ideas and discover if these ideas have merit, and if they do, the effort required to realize them. The challenge is allowing users to leverage as much of the existing infrastructure (so as not to reinvent the wheel), but not be constrained by only what is in production and other guidelines required to meet the production objectives of IT. Rather than just providing a means for users to avoid production guidelines, this process should serve as a mechanism for the business to validate ideas quickly and eliminate initiatives that are ultimately not worth moving to production. This process, hereafter referred to as the "Discovery Process", uses data, tools, business objectives and success criteria as inputs. Upon completion, the outcome can consist of one or more of the following; finalized requirements, working prototypes, deployment estimates, conceptual designs, and business case validation (or invalidation). Figure 3 represents this model.

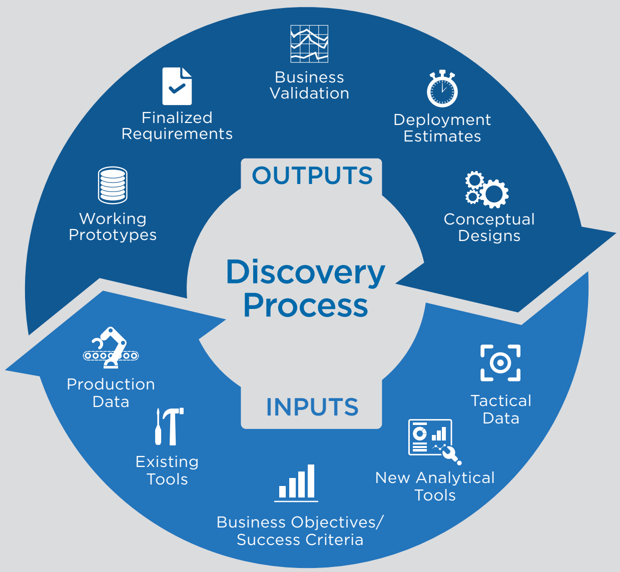

Figure 3—Business Discovery Process

Each discovery process is based upon a preconceived notion by the business where an existing idea has organizational merit. Using this process as a method to test these hypothesis, the user is provided a level of flexibility without being completely dependent on IT. Subsequently, the value of this process is twofold: 1) It can determine if an initiative is worthwhile by proving the business case, and 2) It can define a prototype and/or detailed design such that IT has a clear understanding of what they need to build/deploy as a production-level solution. Done right, this provides a win-win for both organizations.

Because of the level of testing and evaluation, users can define more than high level business requirements. They are also able to define the following:

- **Functional Requirements**—the application functions used to support process specific goals of the business requirements,
- **Data Requirements**—the data components and or integration methods required to support one or more specific functional requirements, and
- **Technical Requirements**—the technologies used to source, create, manage, deliver, clean, and categorize data components that drive the business solution.

These three categories should derive their guiding principles from the strategic Automated Intelligence framework where each strategy's architectural components align with the requirements groupings above. This can be visualized as follows:

As alluded to, in Figure 4, the discovery process does not exist independently, but rather needs to be part of an overall governance process whereby these outputs serve as inputs to determine if these solutions are worth pursuing. If these are worth pursuing, this step will help evaluate the effort, priority, and even a conceptual design that needs to be deployed. A key aspect to enabling this process and making it part of a well-defined and enabling governance process, is to ensure that users have a robust and flexible data environment from which to make discoveries.

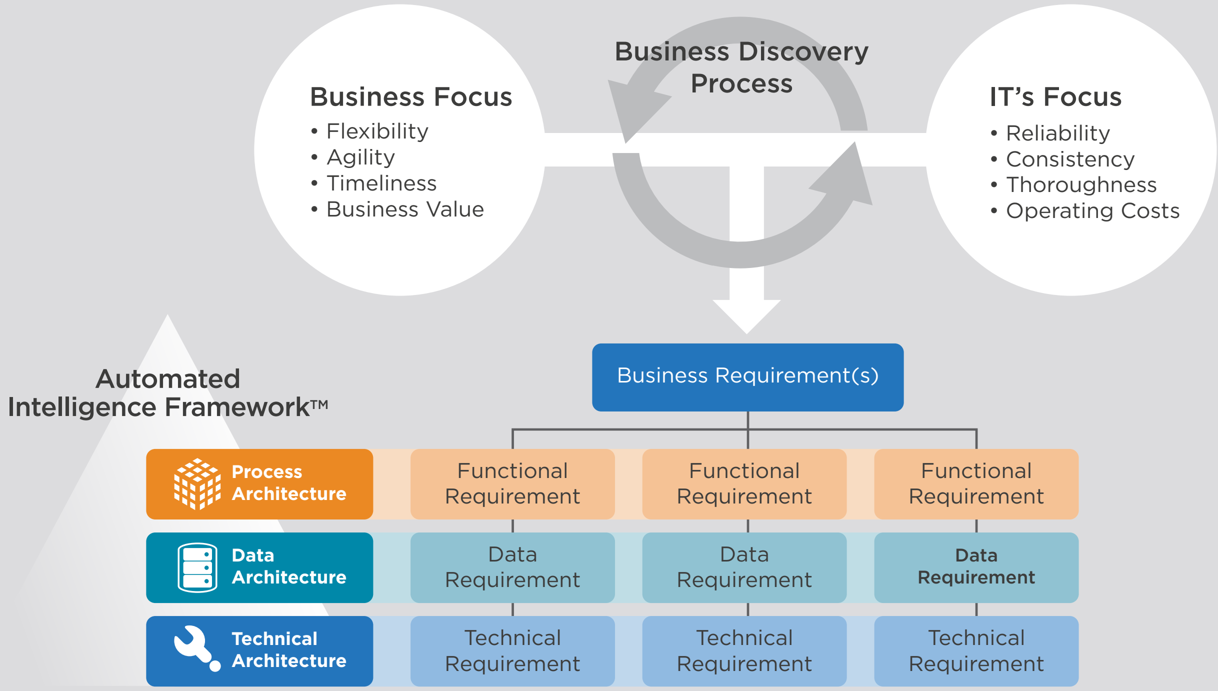

Figure 4—Requirements Hierarchy

Step 3—Make an Enterprise Data Lab an Integral Part of the DMSA Strategy

It is alluded to in step #2, but a key component not represented is a discovery process enabling data environment—hereafter referred to as a Data Lab. The components of the Data Lab are like those found in the Data Management Solutions for Analytics (DMSA) platforms, except for the Operational Data Warehouse (which does not support R&D-based activities). Like the Production environment, the Data Lab may be a single environment or integrated multiple environments with some mechanism to share data seamlessly through user on-demand capabilities.

The Data Lab will need access to some historical and referential data in the traditional data warehouse as a key component to the data environment. Not only access, but the user must be able to create objects in an independent database without the intervention of IT staff. Multiple types of analytics need to be sorted so some components of the context independent data warehouse are required (assuming a traditional data warehouse cannot support all analytical use cases). Finally, users may need access to non-traditional data sources (digital interactions, audio and video files, machine logs, etc.), as well as deep history that may not be available on the traditional data warehouse platform.

Therefore, elements of the Logical Data Warehouse need to be part of the Data Lab environment. Looking at this gives us a subset of the production DMSA model. Note: not all environments are necessarily included in any specific Data Lab group, but these should at least be part of the strategy when setting it up for users and be a part of some future path to enabling the environments not included but will probably come to fruition at some point.

Adding the Data Lab to our Discovery Process enables users to have a specialized, independent R&D data environment from which to run Discovery Process scenarios and build out prototype solutions. As such, this Data Lab (which can be broken into individual environments for different user needs) serves to evaluate new data sources, create new data outputs, and/or test data integration methods. Therefore, by adding this into the mix, we have an extended view of the overall process.

In this process, the Production DMSA environments are used to source some of the data used for the discovery process. Other tactical data is available via some self-service tool(s) (which may include taking production data and transforming it differently). This data will be made available in the Data Lab group so that users can execute via the discovery process. Upon completion of the discovery process, some number of analytical outcomes are evaluated to determine if they belong in a production environment.

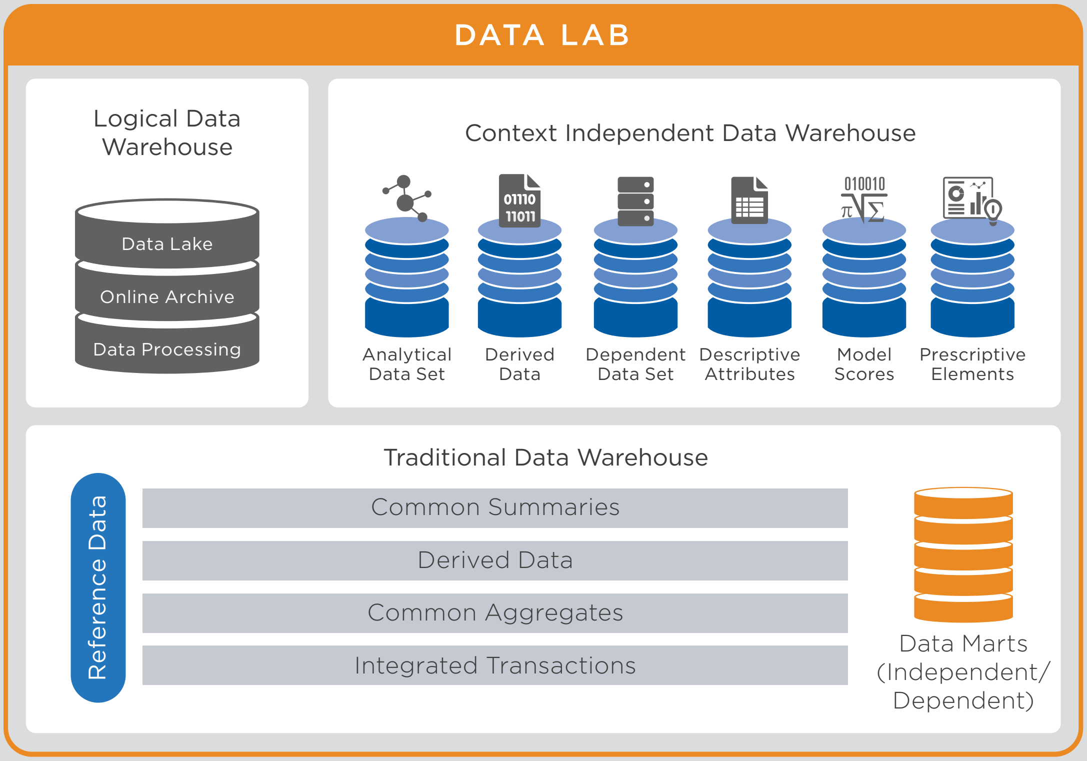

Figure 5—Data Lab DMSA Model

A governance process is in place to evaluate the value and effort of this deployment and prioritize based upon these assessments and the current deployment schedule. If necessary, new data sources and/or data transformations are also deployed to support the new analytical solution. Ultimately, this process (or some facsimile thereof), is intended to allow the business to move quickly on new ideas and come to a point where their initial inferences are validated or invalidated in terms of business benefits. Continuing to manage production processes is a business distraction and building things that ultimately do not provide sustained business benefits are wasted efforts and demoralizing for IT. A process like this (using a Data Lab environment), should alleviate both obstacles.

Some solutions are complex (i.e. they require significant development and potentially new data sources) and some are less involved. Although the latter may not be as impactful, the minimal effort may make these a priority. In either case, not everything can be deployed at once so there should be a well-defined blueprint to outline what is going to be deployed, and when, so that users and IT alike can better manage expectations as they look to enhance overall analytical capabilities.

Step 4—Create a 12–18 Month Analytical Roadmap to Automate Intelligence Within The Organization… and Then Revamp The Roadmap Every Six Months

Innovation is contagious…and by developing a process where innovation is ubiquitous and a framework is in place to guide how analytics are to be developed, applied, and measured, an organization will be well-positioned to succeed. The missing component in this equation is a robust and living roadmap to guide and prioritize the efforts so there is a collective understanding between the business and IT on where things are headed.

Using a nautical analogy, one can have a well-trained and able-bodied crew to manage the ship. The ship can be well-built and in top working order. However, there needs to be a clear roadmap on where the ship is headed to reach the desired destination. Additionally, there needs to be some flexibility in the roadmap to foresee future changes or obstacles and make the necessary adjustments to account for them. This outlines the importance of having an analytical roadmap to support the drive to innovate with analytics.

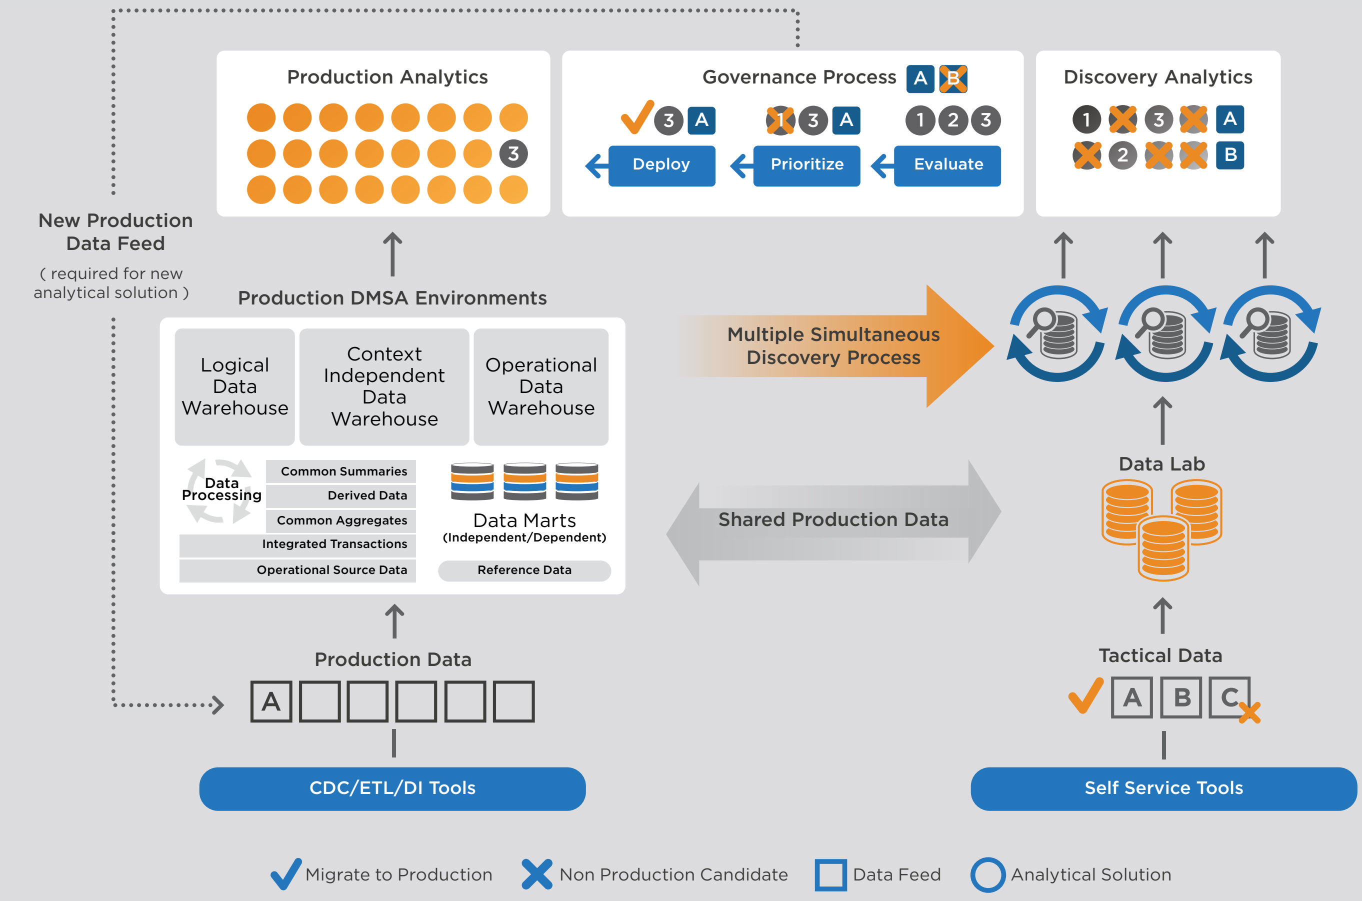

Figure 6—Data Lab Supported Governance Process

A documented roadmap can vary in substance but should adhere to the following basic tenants: 1) Clearly articulate the vision of where you want to be with direct alignment to the enterprise strategy, expected business value, and initiatives underway, 2) Identify where you are today in relation to that vision, and 3) Provide a pragmatic and phased approach to get from where you are to where you want to be.

The vision is defined by an analytical mission (i.e. why the organization is looking to apply analytics) that clearly shows enterprise dependence on the analytics practice to achieve near-term objectives and an already established analytics framework that directly supports enterprise initiatives already underway or on the enterprise roadmap. These should be well-established and provide a series of ongoing guiding principles that drive future initiative requirements.

The gap analysis will take the current list of business initiatives and evaluate them based upon known requirements, business value and juxtapose those against organizational readiness (e.g. process, data, technology). The outcome is a full assessment and current list of initiatives based upon value and organization readiness.

Finally, the assessment information is used to develop a pragmatic roadmap that considers assessment results, related initiatives, and effort to define a phased approach for each of the work streams. Finally, develop a set of projects with related scope, deliverables, and resources required to execute on the defined roadmap.

This roadmap is a great tool to manage expectations and plan human and non-human resource allocation. Additionally, it provides the ongoing narrative about how an analytically-driven practice supports enterprise outcomes. However, it should be a living and evolving blueprint that is flexible enough to accommodate unexpected business changes, tactical additions, and shifts in available personnel and vendors. This belies the need for management oversight that can make these decisions without violating the tenets of the overall vision.

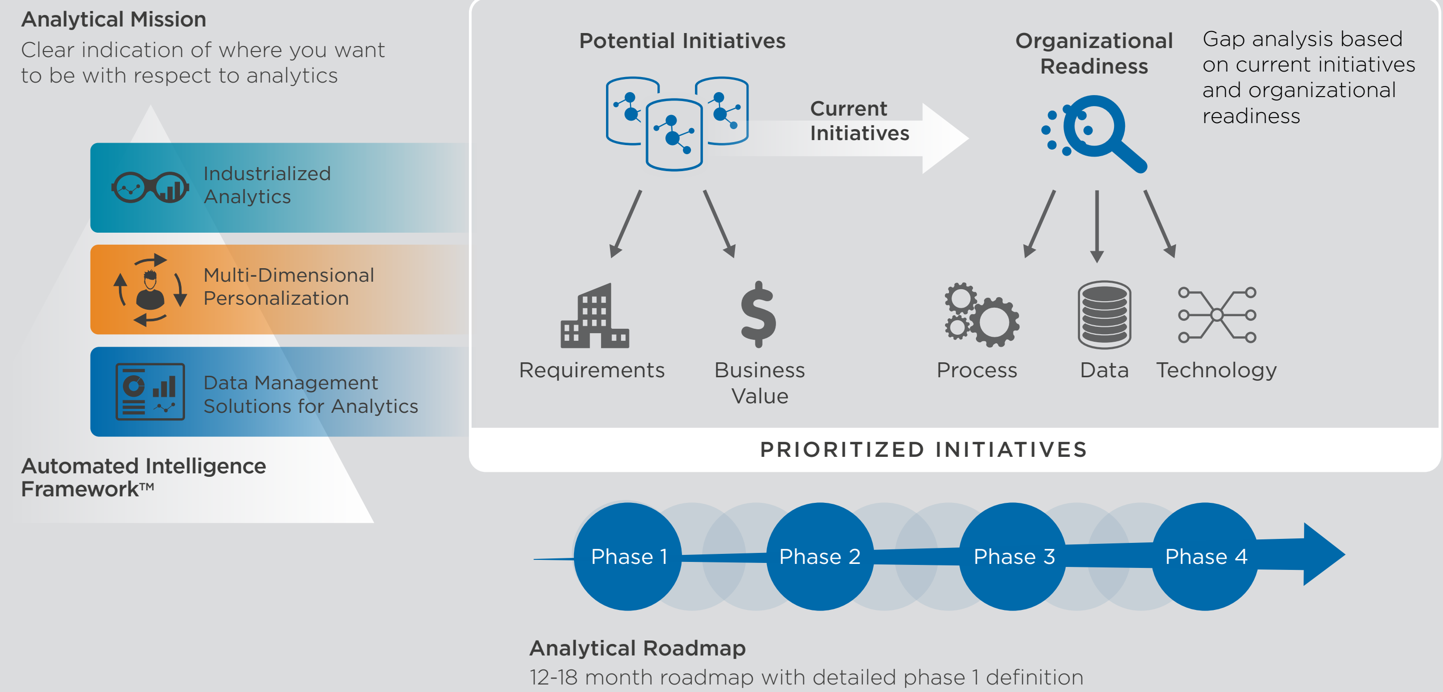

Figure 7—The Analytical Roadmap

Step 5—Institute Actionable Governance: Establishing a Release Schedule Mentality

The discovery process, supporting data labs, and analytical roadmap are intended to give the business a level of autonomy in evaluating potential business solutions and driving enterprise commitment in supporting the activities of the analytical practice. In instances where analytics is a central function in the organization, and especially where a practice has been in operation for some time, multiple simultaneous discovery processes may be running to determine if any number of initiatives are worthwhile to pursue. At most organizations, some type of governance process is put in place to help manage these scenarios, specifically to deliberately and regularly identify those activities that are driving enterprise outcomes and, within those, the activities that are delivering the most value. With that as a charter, governance processes should be seen as imperative and a welcome sight for executives and managers. Many times, however, a governance process is seen (and for good reason) as a mechanism to block the business from moving forward with new ideas.

Governance can be a good idea but the focus must be less on dictating and forcing the adherence to strict guidelines and more on providing the framework to orient the enterprise, provide regular insight into the organization's ability to execute new initiatives, and to identify opportunities to deliver better, faster, and less costly analytically-driven solutions. Ultimately, this minimizes the risk of the business entities building shadow IT organizations to get things done. At its core, the governance process must leverage effort and business value to prioritize initiatives. Additionally, governance should ensure that every existing production component is leveraged fully before new capabilities are deployed (to minimize redundancy). The bigger challenge is determining when one opportunity should be prioritized over another that is already scheduled for release. Another challenge is looking at more tactical opportunities that are relatively simple to deploy but do not necessarily provide strategic benefits (although there are benefits to be gained). This belies the release schedule mentality where you oversee your releases in manageable chunks but make a distinction between enhancements (so the business gets continuous value from things that are not necessarily difficult to deploy but provide significant value) and new development efforts (which are major deployments). The following provides a graphical view of these integrated, end-to-end processes.

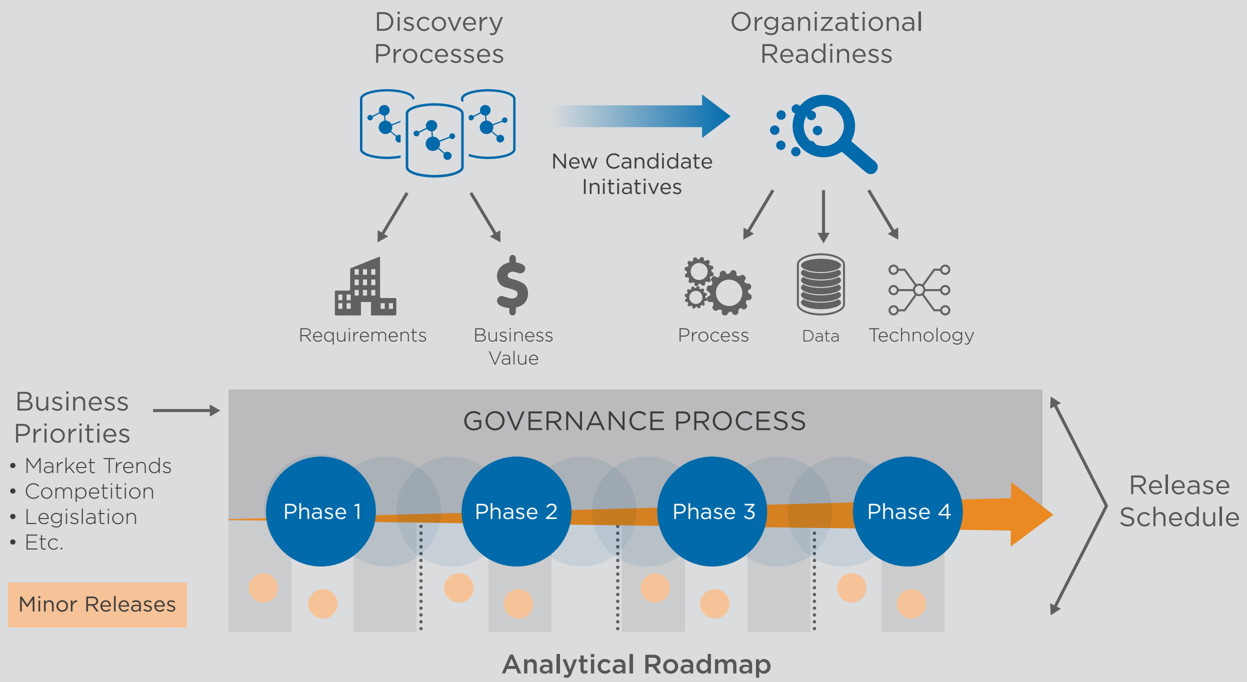

Figure 8—The Governance Model

In this example, business discovery processes produce several candidate solutions and via the governance process these solutions are prioritized and categorized as either minor or major enhancements—subsequently scheduled based on the existing analytical roadmap (which should have outlined multiple phases already). Additionally, the organization will also incorporate changing business priorities as additional input to the Governance Process (which may have changed since the most recent roadmap was defined). The outcome of this is an actionable partnership between IT and the business so that solid, reasoned decisions can drive the definition and potential changes to the release plan over time. It is only through this type of partnership that the goal of operationalizing analytics can truly be met.

The Ultimate Goal: Developing a Culture of Continuous Innovation Through Collaboration

Taking everything discussed into consideration we address the initial disconnect between the business and IT, in that we now have a platform that addresses the business' focus of time to market without compromising the long-term focus of IT on production dependability. The two sets of objectives remain intact because the inherent discovery process provides benefits to both groups of constituents. The Operational Analytics framework is the fulcrum from which the platform remains stable.

Finally, the framework is supported by a robust yet flexible roadmap that provides a guiding light for all associated efforts while managing expectations of all stakeholders and is actively managed with a pragmatic and results-focused governance model. Starting with the original graphic in this section, we might now see something akin to the following:

The outcome of this is an actionable partnership between IT and the business so that solid, reasoned decisions can drive the definition and potential changes to the release plan over time.

Achieving this goal and all those associated with this document will not be an easy task. However, like the model espoused on these pages, users have a framework from which to move away from vendor ambiguity and marketing hype and move towards a pragmatic approach for deriving value from their analytical investments. All the components contained within this document are demonstrable and have shown to be workable solutions in real-world engagements. Instead of touting them as analytical "best practices", the intention here was to provide the tools and techniques to help organizations "practice at being the best".

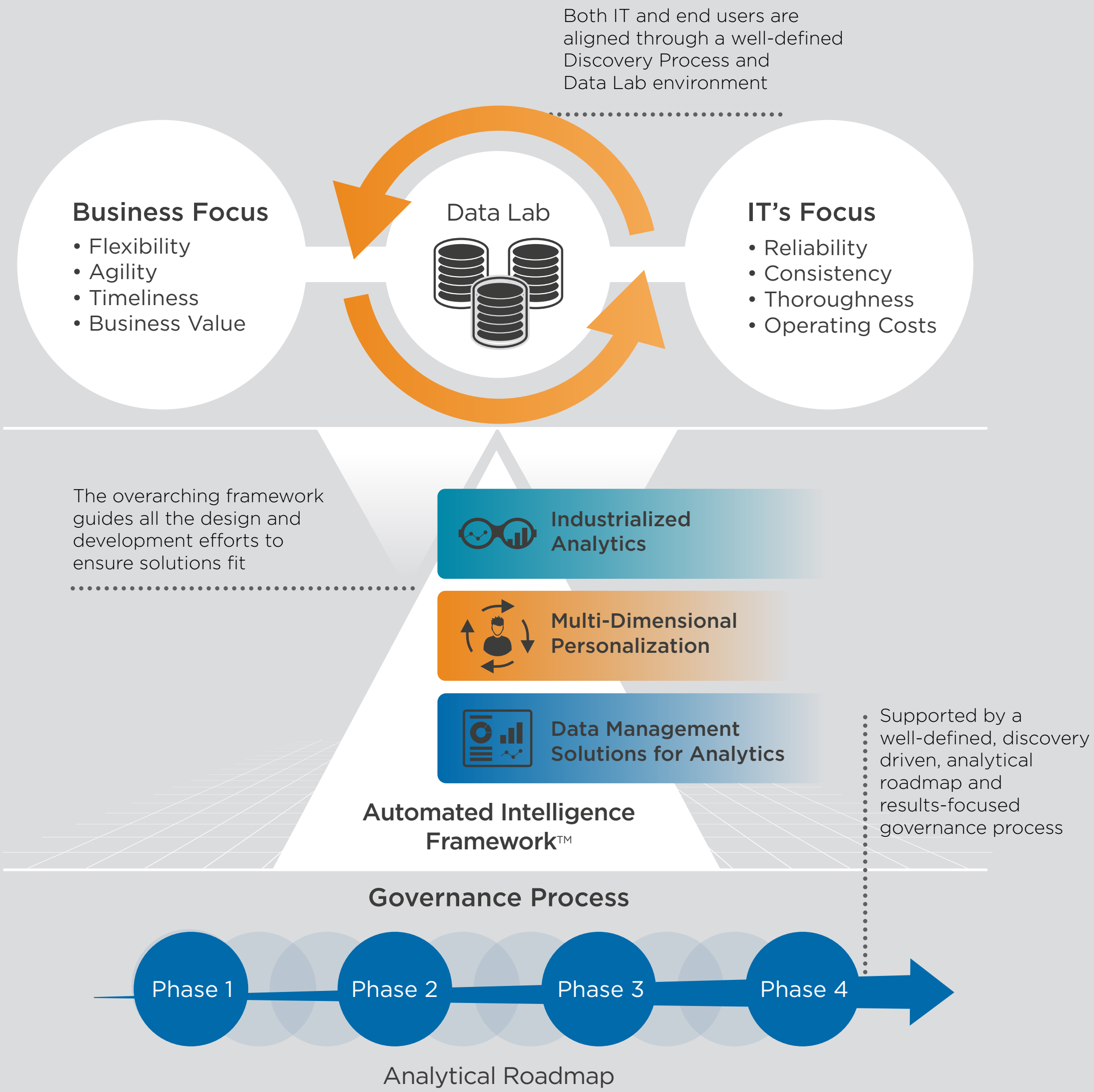

Figure 9—Bridging Business and IT Organizations

This document is part of the Automating Intelligence series. For additional details on the framework outlined and the concepts required for success, please refer to the following additional papers:

- ***Automating Intelligence:** Industrializing Analytics at Enterprise Scale*
- ***Automating Intelligence:** The Real Facts About Artificial Intelligence*
- ***Automating Intelligence:** Taking a Multi-Dimensional Approach to Personalization*
- ***Automating Intelligence:** Recommended Strategies for Applying Recommendations*
- ***Automating Intelligence:** A Pragmatic Approach to Data Management Solutions for Analytics*
- ***Automating Intelligence:** Quality Data, the Cornerstone of Effective Analytics*
- ***Automating Intelligence:** Developing an Organizational Culture to Maximize your Analytical Initiatives*
- ***Automating Intelligence:** The Business-Led Strategy and Framework to Operationalize your Analytical Initiative*

10000 Innovation Drive, Dayton, OH 45342 **Teradata.com**

06.18 EB10128