

TDWI

MONOGRAPH SERIES

JULY 2008

Bridging the Divide: Aligning Analytical Modelers and IT Administrators

By Wayne Eckerson
Director, TDWI Research
The Data Warehousing Institute

SPONSORED BY

TERADATA
Raising Intelligence

tdwi
THE DATA WAREHOUSING INSTITUTE

Table of Contents

Introduction	3
Cultural Divide	3
Perspectives.....	3
Culture Clash.....	4
Where Do They Clash?	5
What's the Upshot?.....	6
Bridging the Divide	7
1. Find a Liaison.....	7
2. Foster Dialogue	8
3. Compromise	9
4. Training	10
5. Implement Sandboxes	11
Recommendations	13



About the Author

Wayne Eckerson is the director of TDWI Research at The Data Warehousing Institute. Eckerson is an industry analyst and the author of *Performance Dashboards: Measuring, Monitoring, and Managing Your Business* (John Wiley & Sons, 2005). He can be reached at weckerson@tdwi.org.

About TDWI Research

TDWI Research provides research and advice for BI professionals worldwide. TDWI Research focuses exclusively on BI/DW issues and teams up with industry practitioners to deliver both broad and deep understanding of the business and technical issues surrounding the deployment of business intelligence and data warehousing solutions. TDWI Research offers reports, commentary, and inquiry services via a worldwide Membership program and provides custom research, benchmarking, and strategic planning services to user and vendor organizations.

About our Sponsor

Teradata Corporation is the acknowledged global leader in data warehouse innovation and analytical solution development. Every day we raise our customers' intelligence to higher levels, making them more focused and competitive by gathering enterprise information and extracting actionable insight.

Teradata elevates enterprise intelligence by giving every decision maker the insight required for smarter, faster decisions. We add value and reveal opportunity across more dimensions than any competing solution.

In every industry and geography, our technologies and expertise make the difference. Simply put, Teradata solutions make companies smarter and give them the competitive advantage to win. To learn more, go to teradata.com.

Introduction

Simmering below the placid workplace environment in companies pursuing data-driven strategies is an internecine feud.

Business analysts who use statistics and sophisticated machine learning techniques to coax hidden patterns and relationships out of large data sets to solve business problems are engaged in a never-ending battle with IT professionals charged with safeguarding corporate data warehouses and ensuring reliable operations of operational applications and systems.

Each side has dug in for a long siege, and neither has contemplated waving the white flag of surrender. Only a few have contemplated a truce, and fewer still have entered into peace talks to end the data wars. The irony is that both sides have much in common. With the right leadership, training, and technology, these antagonists can find common ground and resolve their differences. Each side can get what it needs, if not what it wants, providing their organizations with heightened productivity and bottom-line benefits.

This monograph examines the cultural divide between analytical modelers and IT administrators. Anyone who has managed or mediated between these two groups knows that there is a tremendous gulf that exists between them—organizationally, culturally, and functionally. One IT professional I spoke with said, only half jokingly, “You can even identify individuals in each group just by looking at them.”

This report analyzes the perspectives and prejudices each group has about the other, examines the cultures that give rise to these prejudices, and assesses the business impact of the mutual antagonism. Finally, it describes techniques used by business and technical managers to close the cultural divide and get the two groups working together effectively to deliver optimal results for their organizations.

Cultural Divide

Perspectives

How Analytical Modelers View IT. Analytical modelers view IT as gatekeepers to the data and bottlenecks to getting things done. At best, they find IT unresponsive to their requests for new sets of data, and at worst, demeaning. One modeler told me, “If you’re not in IT, they think you are technically incompetent and won’t give you the time of day.” When IT does accede to their requests, analysts find them incredibly slow in getting things done. “I can do a lot of things really fast—it takes me two weeks to do things that take IT six months.”

To analytical modelers, IT seems paralyzed by process and blinded by methodology. IT’s systematic approach to development seems more like a smokescreen to deflect criticism, should a project fail, rather than a way to ensure a successful outcome. At a public forum this year, I listened to a manager of an analytical modeling group rant for an hour (with some gentle goading) about the myriad shortcomings of IT and data warehousing managers. His basic message was that IT doesn’t understand business processes and therefore can’t be relied upon to deliver accurate, relevant data in a timely manner. As a result, his group gets its own data and doesn’t rely on the IT department, which he thought could be eliminated to save the company money.

In short, modelers see IT as a big obstacle to be circumvented, and that’s putting it nicely.

IT's Perspective of Analytical Modelers. The IT department, on the other hand, pulls its hair out dealing with analytical modelers. They view them as fickle pickles—people who keep changing their minds about what data they need. Modelers are the users who deliver the classic data warehousing line when asked what data they need: “I don’t know. Just give me all of it.”

IT also sees analytical modelers as security breaches in the making, since most of them want to download large volumes of data, including customer records, to their desktops or laptops, which can easily walk out the door or be stolen. IT also questions modelers’ ability to create properly formed SQL and reconcile data of dissimilar grain for analysis; they suspect modelers of creating poor-quality data sets. They know analysts create the majority of the analytic silos throughout the company running on Excel, Access, and other databases, which are IT’s job to consolidate and eliminate.

IT also views analysts as CPU hogs and performance sinkholes who issue runaway queries and kick off CPU-intensive routines that erode performance on database and application servers. The resulting poor performance causes users to complain and blame IT for the problems. Finally, IT fears that analytical modelers who rummage roughshod through corporate databases will inevitably crash a system. Of course, IT will get blamed for the failure, not the analysts, and IT will have to work overtime to fix the problem.

So, in short, IT views analytical modelers as a threat to the integrity of the corporate data upon which their reputations, careers, and chances for a normal work schedule hinge.

Culture Clash

What is the root cause of these perspectives, or should we say prejudices? On the surface, the cause is a simple one. These two groups are part of two distinct cultures: the business culture and IT culture. Because of economies of scale, most companies segregate IT from the business and place it in a distinct organization with its own jargon, responsibilities, incentives, goals, and career paths. Such artificial separation intrinsically breeds suspicion, distrust, and eventually conflict. (See Figure 1.)

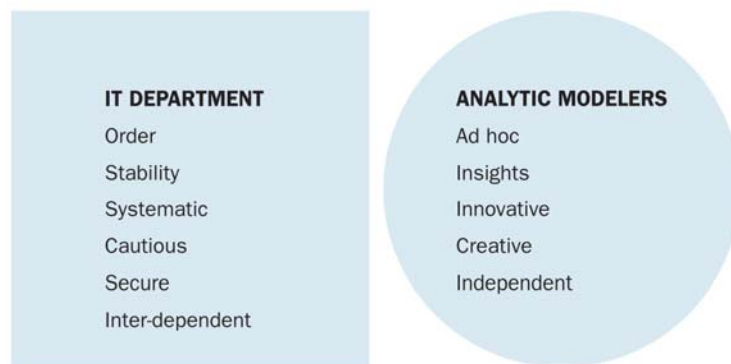


Figure 1: Culture clash. Culturally, IT and analytical modelers are polar opposites, creating tension, distrust, and conflict. One DW professional confided, “The people in these two groups even look and dress differently.”

But a deeper cause comes from their job responsibilities and incentive and performance plans, which put them in direct conflict.

The IT Imperative. The job of the IT department is to keep systems running around the clock, including data warehouses, many of which are now mission-critical systems. This requires IT to manage complex processes and meet stringent system level agreements for uptime, performance, and data quality. IT's job is to maintain data and application security and comply with regulations regarding privacy, controls, and access.

Thus, it's not surprising that the IT culture revolves around order and stability. IT personnel must function as systems police, ensuring the safety and security of organizational systems, data, and applications that run the business on a daily basis. They must take a comprehensive and systematic approach to development to ensure that changes to one system don't adversely affect any applications and users downstream.

In short, IT is under a lot of pressure to keep things from breaking. Like linemen in American football, they get little visibility and recognition unless they make a mistake.

The Business Imperative. In contrast, the business side of the house carries a different set of goals and responsibilities. To achieve success in an ever-changing and competitive marketplace, analytical modelers work in an ad hoc fashion to deliver insights potentially worth millions of dollars in revenue. They work iteratively, rather than systematically, constantly experimenting with new ways to tune their models. As one analytical modeler told me, "Our prerogative is to push the envelope on the business side and produce actionable information."

Unlike the IT department, analytical modelers are like football running backs and receivers, who get all the attention and praise for solving high-profile business problems.

Where Do They Clash?

So where specifically do these two groups clash? There are many incidents that trigger conflict, but here are three of the most common.

1. Runaway Queries. To create models, analysts search through corporate databases, looking for the best data and variables with which to create their models. This often involves using BI products as glorified ETL tools to download large volumes of data to an analytical workbench, desktop database, or server. Of course, DBAs see runaway queries as a performance nightmare and do everything they can to halt or reschedule these queries to avoid slowing query performance for other users of the data warehouse. They also view the data dumps as potential security breaches.

2. Analytic Data Sets. Once modelers acquire various data sets, they then merge, manipulate, and aggregate the data into a single wide table—or analytic data set—against which they run various algorithms to create a model. They also need to regularly recreate these data sets to update their models with the latest data. Meanwhile, IT departments view these data sets as silos of inconsistent data that is their job to stamp out, and they question the SQL skills of the analysts to accurately query and integrate disparate data. As a result, the IT department is reluctant either to create the analytic data sets or give modelers free run of the database.

Unfortunately, without strong cooperation from IT, analysts can struggle to create or update analytic data sets in a timely manner, which affects the accuracy and performance of their models. Essentially, the business ends up making today's decisions with last month's data.

3. Scoring Processes. Third, to deploy the models, analysts may download all relevant records to their desktop, score them, and upload them to the data warehouse, which floods the corporate network and may cause synchronization problems with the database. If they need to score millions of records, the modelers will upload procedures to the database and score the records there. Since scoring takes an incredible amount of processing power and time, there's a good chance this step can quickly bring a database to its knees or even corrupt it if the code is not written properly.

What's the Upshot?

So when IT and analysts clash, what's the upshot for the business?

It's simple: IT doesn't want to get burned, so it slams shut the data gate, denying analysts complete and unfettered access to the data, which they need to do their jobs. In turn, the analysts, who are generally clever, independent-minded, and ambitious, find a way to circumvent IT and get what they want. More specifically, this cultural divide has three major impacts upon the business:

1. Timeliness Suffers. The first impact is that analysts don't deliver models in a timely fashion. Analysts spend upwards of 50% of their time accessing, exploring, and manipulating data—tasks that the IT department could assist with. TDWI Research shows that analytic modelers spend 43% of their time exploring and preparing the data and 23% creating the model, which itself requires iterating back through exploration and data preparation task.¹

"When I came on board, it took us four months to create a model and four months to deploy it," said the vice president of risk management at a major financial services company, who set a goal to deliver new models every month, which would help the company generate millions in additional revenue.

2. Model Quality Suffers. Second, model quality suffers when analysts and IT aren't aligned. When denied open and easy access to data, analysts perform fewer iterations of their model because IT has put up too many roadblocks. As a result, their models are less optimized, resulting in lost revenue for the business.

3. Inefficiency and Waste. Finally, analytical modelers who spend 50%+ of their time managing data are not performing the work they were hired to do, which is to create and analyze models. Instead, they spend half their time working as glorified IT administrators. Freeing analysts from having to prepare and manage data enables companies to radically improve the productivity of their analysts.

¹ The remaining time is spent defining a project (13%), scoring records (12%), and managing models (9%). From the TDWI Research Best Practices Report *Predictive Analytics: Extending the Value of Your Data Warehousing Investment*, Wayne Eckerson, Q1 2007. Available free at www.tdwi.org/research.

Bridging the Divide

Many people recognize the need to align IT and business analysts, but few succeed in doing so. That's largely because most of the solutions for bridging the divide are organizational and cultural in nature. In other words, the solution lies in changing attitudes and behaviors, and this type of change doesn't come easily or quickly.

From my discussions with various business and IT managers, I've discovered five ways to bridge the divide between IT and analytical modelers: 1) Find a liaison; 2) foster dialogue; 3) compromise; 4) train; and 5) build analytical sandboxes. (See Figure 2.) Let's go through each of these.



Figure 2: Keys to bridging the divide. There are five ways to bridge the divide between analytical modelers and the IT department.

1. Find a Liaison

Every organization that has aligned analytical modelers and IT administrators has had a strong leader acting as a liaison between the two groups. Liaisons are people who straddle both worlds—they talk the language of both business and IT and manage to align the two groups through force of personality and partnerships.

“I’m one of the few who stands in the middle and tries to make sure people work together,” says a vice president of strategic technology at a major credit card and financial services firm. “I’ve established strong partnerships with people in the business, but neither side of the fence sees the other’s position. They don’t speak the same language.”

Liaisons are generally senior-level professionals who have a long history in the business or their industry. They have a lot of energy and never miss an opportunity to evangelize the need to establish stronger ties between the two groups. There is no pattern to where they work or to whom they report. Sometimes, the person resides in the IT department and knows how to reach out to business partners, and other times, the person works in the business but is intimately familiar with the ways of IT.

“I am the truest hybrid you could meet,” says a vice president at a financial services firm. “I’ve worked on the business side all of my career but focused on executing IT initiatives. Thus, I can

talk to the IT guys and brainstorm with them how to solve a problem. I can say to them, ‘This is what these guys want, and let’s do it this way.’”

Despite the critical role they play, few companies assign people to act as liaisons. Most liaisons adopt the role voluntarily because they see the need and want to help the company succeed.

Organize the Analysts. Of course, having a liaison implies that there are two groups with which to liaise. But often, that is not the case.

In many organizations, analytical modelers—as well as run-of-the-mill business analysts and power users—are scattered throughout the company, isolated in individual departments. The problem with burying your analysts in departmental groups is that they don’t feel a strong allegiance to the company. Without a collegial atmosphere and a clear career track, many are tempted to pursue higher-priced offers at other companies.

To rectify this problem, some companies have moved all their analysts—including power users (i.e., report developers), financial analysts, and analytical modelers—into a single business unit, such as finance, marketing, or IT. The analysts retain dotted-line responsibility to an individual department and spend many hours in meetings there. But all the analysts are physically located in the same office, report to the same manager, and are paid by the same group. “Our rationale was to foster camaraderie and career growth for the analysts and make it easier to get them the resources they need to perform their jobs at the highest level,” said one executive who uses this approach.

2. Foster Dialogue

The second way to bridge the divide is to foster dialogue between the two groups. Ongoing dialogue increases mutual understanding and respect, and softens prejudices and knee-jerk reactions. A consultant who specializes in helping analytical modeling teams exploit analytical extensions said, “I feel like a marriage counselor. I try to promote a dialogue between the two groups.”

There are several ways to foster dialogue among warring groups.

1. Socialize. Get the two groups to socialize with each other outside of work. One manager takes both teams to baseball games or bowling on a periodic basis. Besides promoting mutual understanding, social interaction reduces the possibility of open hostility. As the manager says, “It’s harder to get mad at someone whom you know socially and have become friends with.”

2. Clear-the-air meetings. A more formal way to boost mutual respect and understanding is to have a joint meeting where the two groups outline their frustrations with each other and then identify things they have in common. Once the groups recognize their common ground—that they are working for the same company and toward the same ultimate goals—they find more reasons to cooperate and avoid conflict.

3. Formal Dialogue. Finally, it’s important to set up formal channels for regular, ongoing communications. For example, an online retailer holds two sets of biweekly meetings. One meeting is for executives of the two groups—specifically the VP of analytics and business intelligence, the CIO, and the head of data warehousing. The other is for managers of the two

groups. The executives look at longer-term, strategic issues, while the managers focus on tactical concerns involving projects, query performance, and reports, among other things.

“It should never be a surprise to IT executives what analysts are working on that could adversely affect their systems,” says the vice president of analytics and business intelligence.

Empathy. Ultimately, the purpose of all this dialogue is to engender empathy on both sides of the aisle. For example, an analytical modeler, while frustrated at the slow pace of IT, now understands why they act the way they do:

IT’s main concern is not having applications in the data warehouse go down so their people don’t have to work nights and weekends. Business people don’t always recognize the complexity of keeping those systems working. And IT [staff] don’t get a lot of recognition or respect and aren’t involved in business decisions that affect them, and that’s hard for them.

3. Compromise

Once the two sides understand each other and have developed a degree of empathy for the pressures and responsibilities each faces, then it’s time for both groups to roll up their sleeves and change attitudes and work habits.

In some cases, this means analysts need to learn the language of IT, which is SQL. It also means that once given access to the corporate jewels (i.e., the data warehouse), the analysts should adopt some IT processes, such as modeling rules and naming conventions, especially if they want IT to deploy their models in a quick and efficient way. They also need to realize that, despite advances in technology, which we’ll discuss shortly, they need to be cognizant of other applications running on the system and adjust accordingly. That might mean running a sizable query at night instead of on a weekend to avoid a scheduled backup or unplanned maintenance work.

On the IT side, compromise means giving analysts access to data warehouse data and giving them space to add, delete, and manipulate tables and data as they see fit with the required systems resources (e.g., storage capacity, CPU threads, and memory) to get the job done. Sometimes database administrators restrict access reflexively—either because of timeworn habits developed when storage, memory, and CPUs were expensive resources, or because they are afraid of what analysts might do with such unfettered access.

Compromise might also mean loading the data warehouse with sufficient attributes and history so analysts aren’t tempted to circumvent the data warehouse and extract data directly from operational systems. IT can also precalculate aggregations and transformations that analysts commonly apply to the data when creating models. For example, analysts may aggregate total customer spending by day and by week to optimize their models.

Bottlenecks. The financial services firm mentioned earlier (that required eight months to deploy a model) discovered three major bottlenecks to faster deployment:

1. Analysts were entirely reliant on IT to create data sets for them.
2. There were not enough system threads available, so jobs were queued, causing long delays in processing models.

3. There were no clear data definitions, making it hard for analysts to really understand the data they received.

Solutions. The company came up with a number of solutions, which forced both analysts and IT to make some compromises.

1. Analysts had to learn SQL so they could access the data directly instead of asking IT for data dumps. They also had to convert their analytic code to SQL to make it easier and quicker for IT to set up a scoring routine and run it.
2. The IT group had to implement a higher-performance data warehousing platform to remove systems bottlenecks that caused performance problems. IT also had to grant analysts direct access to the data warehouse, something it had never done before. To optimize such access, it implemented user-friendly denormalized schema and database views to shield analysts from the complexity of back-end tables and naming conventions.
3. IT loaded all the data that analysts might possibly need, including more than 8,000 attributes and up to 60 months of history for the most commonly requested fields.
4. The business also implemented a governance council to clean up and standardize data definitions.

By taking these actions, the financial services firm achieved its goal of delivering models in one month or less.

4. Training

The fourth way to bridge the divide is training. To make the changes described here requires both groups to pick up new skills. For example, the financial services firm just mentioned implemented several training classes for analysts. One class was on basic SQL; another was on SAS SQL; and a third covered how to create, drop, and manipulate tables in Teradata. They learned which functions were best handled by Teradata or SAS and created their programs to leverage the strengths of each.

Conversely, database administrators need to learn how to configure and tune a mixed workload environment that supports a combination of users (e.g., casual users, power users, analysts, and administrators) running a variety of processes (e.g., loading, updating, transforming, querying, and reporting). This may require hiring consultants to show the database administrators how to configure a mixed workload system in their own environment.

Buddy System. To solidify the lessons learned during training, some companies use a buddy system. Here, a novice analyst is paired with an experienced analyst who already knows SQL and how to manipulate the DW database. This way, novices have someone familiar to call in case they get stuck or forget what they learned during training.

Patience. It was also important to set the expectations of the IT department and get them ready for an onslaught of poorly formed queries. One manager said, “I sat down the DBAs and said, ‘Look, you are going to get some of the worst queries you’ll ever see. Remember they are analysts, not technologists. You are not to yell at anybody! Give them a call, cancel their query,

and say, ‘How can I help you out? What are you trying to do?’ And they did and were very supportive.”

5. Implement Sandboxes

The last way to bridge the divide is to create analytical sandboxes. These are dedicated environments that permit analysts to do all the things analysts love to do—such as creating tables, loading data, merging and aggregating, combining fields, and running queries—without impinging on other users of the data warehousing environment. In many respects, analytical sandboxes are the key to making the other four practices work, because the technology gives each group what it needs without making undue compromises.

eBay. For example, eBay recently implemented analytic sandboxes, which it calls Prototyping Environment for Teradata or PET, to better meet the needs of its analysts and eliminate their need to create renegade data marts.

Oliver Ratzesberger, director of information management architecture and operations at eBay, says,

As quickly as we integrated renegade data marts into our data warehouse, new ones popped up. So we created virtual data marts in the enterprise data warehouse and user adoption has gone through the roof! We give users their own dedicated database to which they can add data and combine it with data from the enterprise data warehouse. They use these environments to prototype new analytical solutions. When they feel their solution is production ready, they turn it over to IT for development. Now the business gets the freedom to create the tables they want without having to worry about running a SQL Server box under their desk.

Two Types. There are two ways to set up an analytical sandbox. (See Figure 3.) One way is to create a physically distinct data mart and then load it with all the data analysts might need in a form that suits them—wide flat tables. Analysts can then add external data to these data marts. This is the traditional approach that organizations have used to create sandboxes.

A second way—and the one that eBay supports—is to create virtual sandboxes by partitioning spaces within a data warehousing database. IT will load a predefined analytic data set from the warehouse into the virtual sandbox on a regular basis using ETL tools. Analysts can copy additional data from the warehouse into the sandbox and merge it with external data.

In essence, a virtual sandbox gives analysts free rein to create tables, add data, and merge, manipulate, and calculate data sets as they see fit without impacting performance for other users of the system or jeopardizing the integrity of data warehousing data. Once an analyst finishes building a model, they then throw it over the wall to IT to implement in the production side of the DW database.

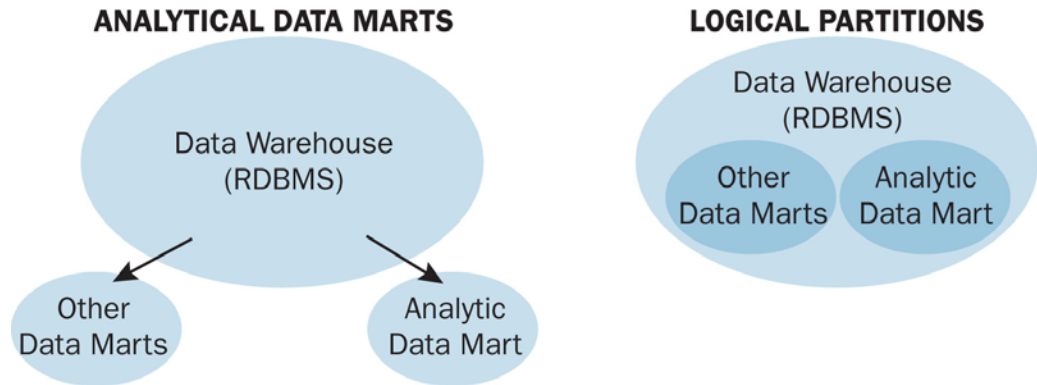


Figure 3: Types of analytical sandboxes. There are two ways to construct an analytical sandbox. The classic way is to create a dependent data mart set up exclusively for use by analytical modelers. A newer way uses database partitions and views to create a logical mart within the data warehousing environment.

Workload Management. Recent advances in database technology—specifically, sophisticated workload management and partitioning schemes—have made the virtual sandbox an increasingly popular option. However, managing this environment can be challenging for DBAs who don’t know how to configure a mixed workload environment or whose database environments aren’t powerful enough to support dozens or hundreds of partitions for analytical modelers. As a result, companies may need to spend money to hire consultants and upgrade their database environments to support in-database analytical sandboxes.

In addition, IT has to set some ground rules for analyst behavior within the sandboxes. While analysts have tremendous freedom, there may be some constraints on how much data they can load and when they can kick off certain types of processes. Obviously, there must be strong communication between the two groups to avoid problems. IT also has to interview analysts to discover the type and scope of data they want automatically loaded into their sandboxes via ETL.

There are trade-offs to using physical and virtual sandboxes. (See Table 1.) A physical sandbox is simpler for the IT department to set up and administer and offers predictable query and processing performance for analysts. However, it is more expensive because it requires installing a separate server environment. It also requires the IT department to create ETL and replication routines to physically move data from one machine to another over a network, which increases the complexity of data migration compared to a virtual sandbox. It also increases data redundancy and proliferates analytical silos.

A virtual sandbox, on the other hand, reduces data redundancy, gives users access to the freshest data possible in the warehouse via database views, and halts the spread of analytical silos. It doesn’t require organizations to invest in a separate server and database environment or duplicate and move data across a network. As such, it offers a simpler architecture and lower costs. However, it does require administrators to implement workload management, which is complex to set up and may require organizations to upgrade their databases and server environments and hire consultants to set up the environments initially.

	Physical Sandbox	Virtual Sandbox
Pros	<ul style="list-style-type: none"> • Simplicity • Predictable performance • No workload management 	<ul style="list-style-type: none"> • Lower costs: Leverages existing hardware, software, and IT resources • Simplified application and data migration • Fresher data via database views • Reduced data redundancy and analytic silos
Cons	<ul style="list-style-type: none"> • Application and data migration challenges • Higher cost: Additional hardware, software and IT resources • Less fresh data that requires replication and ETL • Increased data redundancy and analytic silos 	<ul style="list-style-type: none"> • Variable performance in sandbox • Workload management is complex to set up and tune and may require a database server upgrade and consulting services

Recommendations

In summary, analytical modelers and data warehousing administrators hail from different cultures. Most organizations have institutionalized these cultural differences by segregating the groups into different departments, and the groups have failed to establish formal channels of communication that would enable them interact and engage in a constructive way.

To overcome the cultural and organizational divide and get these groups working together harmoniously for the good of the organization, business and technology managers need to:

- **Find a liaison.** Find a liaison in either the business or IT who knows how to talk with both groups and get them working together in a productive way. Ideally, the person has significant clout in the company and can leverage his or her experience and partnerships to overcome conflict and align the two groups.
- **Foster dialogue.** Foster formal and informal channels of communication to increase mutual understanding between the groups and find common ground to solve problems where both sides emerge as winners.
- **Get the groups to compromise.** Once the groups empathize with each other, get them to commit to making changes that will help meet the fundamental needs of the other group. This is never easy, but can be done with strong leadership.
- **Provide training.** While the spirit is willing, the flesh is weak. Both sides may want to change for the benefit of the organization, but may lack the skills and knowledge to make

it happen. A multi-tiered training and buddy program can help grease the wheels of change.

- **Implement analytical sandboxes.** Finally, the IT department needs to implement analytical sandboxes that give each side what it wants. The analysts get reasonably unfettered access to data, and IT gets to lock analysts in a circumscribed environment where they have license to play without harming other users on the system.

Implementing the recommendations described here will improve the productivity of your analytical modelers and provide significant lift in the accuracy and performance of their models. This, in turn, can generate millions of dollars in additional revenue or cost savings for your organization.