



No matter the client we've serviced, the products we've worked with, or the size of the project we've delivered, every habit continues to drive our teams to the right decisions.

The Seven Habits of Highly Effective BI Teams

As consultants in the software development industry, Quadrus is always on the lookout for techniques and methods to add to our best practice arsenal. To this end, our Business Intelligence (BI) practice has developed seven essential tenets to reinforce those ideas that we have found to be most useful in our BI and data warehouse efforts. We call these tenets, *The Seven Habits of Highly Effective BI Teams*.

Because they are habits rather than specific problem-solving approaches, each can be applied in almost any BI situation regardless of client, industry or methodology. In combination, each fosters a way of thinking – a view of BI development – which helps steer us away from the problems, issues and common gotchas that cause BI projects to fail.

The following is a list of these seven habits along with methods and techniques that encourage their application.

Habit #1 – “Embrace the User and their Business”

Studies suggest that data warehouse failures are not a result of technology or design issues but a consequence of users simply rejecting the system. Numerous explanations and ideas have been offered in the literature as to why this phenomenon exists. Our group has its own theory:

“Data Warehouse/Business Intelligence development teams and software vendors forget that real users and businesses are the true customers.”

This is a somewhat tongue-in-cheek perspective that clearly bears scrutiny, but we believe it often holds true. Consider the following; for years, a major vendor of reporting software offered only minimal support for spreadsheet integration. Though an excellent reporting tool, our clients were often forced to eliminate the vendor from their product assessments due to the lack of this feature.

“We are reminded that our primary goal is to ensure the information we deliver is valuable, the tools work, and the solution makes sense.”

When we asked the vendor why the product did not support spreadsheet integration – and explained that this was often the reason we could not recommend it – we were told that the product was designed to eliminate spreadsheet reporting outside the data warehouse. As such, the vendor told us, “A properly designed data warehouse shouldn’t need this type of integration.”

We think this type of thinking reveals a peculiar lack of sensitivity to the practical aspects of BI within the industry. Although members of our BI practice team are anything but advocates of proliferating spreadsheets, we recognize that tools such as Microsoft Excel offer additional value as extensions to most data warehouse implementations.

We don’t consider Microsoft Excel, Microsoft Access and other desktop products as threats but rather as valuable tools in the user’s potential reporting arsenal. More importantly, we understand that the BI effort will only be successful if the solution fits and integrates effectively with the user’s day-to-day activities. Such activities include the ability to perform analysis through familiar reporting paradigms such as spreadsheets. In other words, we believe that a BI solution must not only work, but it must work in a way that fits with the user’s business activities.

The first habit on our list – *Embrace the User and their Business* – reminds our team to steer away from the thinking exhibited by this vendor; that we must not only focus on what information is needed but also be sensitive to how the user will want to interact with the data. It reinforces us to recognize the specific reporting needs of each user as well as the importance of selecting the right reporting tools based on these needs. It also reminds us that our primary goal is to ensure that the information we deliver is valuable, the tools work, and the solution makes sense.

This habit also encourages our team to determine how the users will work with the new system and whether they are likely to reject it. Investigations include how users will perform their current day-to-day reporting activities, how data will be drawn from the existing reporting environment, and how information will be manipulated in associated processes. Team members are also asked to perform the actual reporting procedures that will be supplemented or replaced by the new BI solution.

Such practices provide our team with hands-on knowledge of the existing reporting environment – its positives and negatives – and gives us a better understanding of how our users will react to the new technology. Once armed with this information, we have found that the team is better equipped to build a successful solution that is widely used and accepted.

Does employing this habit require additional effort? Absolutely! But we feel that this effort significantly outweighs the risks of delivering a system that is ultimately rejected by the users.

“Poor data quality isn’t just the effort involved in fixing data errors or changing business processes. Poor data quality robs a data warehouse of its ability to be trusted, and users who don’t trust a data warehouse won’t use it regardless of why trust has been lost.”

By keeping our projects rooted in the realities of our users’ day-to-day experience, this habit has made our team more sensitive to real needs and – as result – better consultants.

Habit #2 – “Seek Out and Tackle Data Quality Issues”

As most BI practitioners know, the effect of poor data quality isn’t just the effort involved in fixing data errors or changing business processes. Poor data quality robs a data warehouse of its ability to be trusted, and users who don’t trust a data warehouse won’t use it regardless of why trust has been lost. Our number two habit - *Seek Out and Tackle Data Quality Issues* - expresses our belief that source systems must be treated with suspicion, business rules must be tested, and that data quality issues must be managed defensively. The stakes are simply too high not to.

How do we incorporate this habit into a typical BI project? Here are a few examples.

To deal with data quality issues, many BI teams develop simple ETL filter processes that detect only the most basic types of data input errors (Bad Date, Value Not Numeric, Value Out-Of-Range, etc.) Records that don’t pass the filter tests are simply updated in place or rejected outright. We consider this level of detection completely inadequate for a properly designed data warehouse.

The next tier of data quality detection commonly implemented is a collection of filter algorithms based on user-defined business rules. This is a more complete approach to data quality detection and has the potential to provide good quality data. In practice, however, organizational business rules are insufficient to ensure proper data quality management. The problem stems from the fact that business rules are informally maintained and are multiply owned. As such, when these rules are translated into formal filtering and error detection scripts, they often cause ETL failures due to inconsistencies in their application at source. Over time, the annoyance of these failures forces many of these business rule filters to be relaxed to the point where they no longer act as safeguards within the ETL.

To avoid this situation occurring within the ETL processes, our team looks to the second habit. Before we build any ETL processes, we perform Data Quality Audits based on the business rules provided by the users. These audits offer a means to seek out issues before any information is brought into the data warehouse.

To further improve on this process, we also ensure that a data steward from within the business is assigned who is given the responsibility for collecting and reconciling business rules across the organization. This way we eliminate the problem of dealing with the problem of “many truths”.

“One shouldn’t assume that a suite of tools from the same vendor is always the best approach. Vendors who offer the best reporting tools may not stack up well on the data movement side and vice versa.”

As the data quality audits are performed, the data steward is kept up to date with the results to ensure that the business rules are consistent with physical data and processes. Once all of our data quality audits have been performed successfully, we then build the final rules into our ETL scripts.

But we don’t stop there. We then tackle the data quality issue by creating processes and data structures within the data warehouse to monitor and record when a violation has occurred. These methods are designed to integrate with the dimensional models that support the reporting layer within the data warehouse. This allows the user community - and specifically the data steward - to easily review data quality issues as they arise. The end result leads to updated business rules, changes to business processes, a change in the ETL scripts, correction and re-processing of the invalid data, or some combination thereof. Incorporating these steps into the BI development lifecycle ensures that data quality issues are uncovered and dealt with long before they have the capacity to impact the success of the final product.

Habit #3 – “Research, Test and Pilot Technology”

As part of our BI consulting work, we are asked to perform assessments of corporate BI/DW projects that have run aground. What we often find is a project with an excellent ETL design, exceptional information architecture, and highly skilled teams, which failed simply because the reporting tools didn’t fit or couldn’t be made to work. When we discover this situation, we ask to see the results of the project tool assessment phase. As often as not, this request returns blank stares.

When data warehousing was still in its relative infancy, BI projects were often derailed by technology. The tools were unproven and the BI waters were largely uncharted. One had no choice but to jump in and hope for the best. That time has come and gone, and today there really is no excuse for a BI project to fail because the wrong tools were selected. There are dozens of data integration and reporting technologies available and it is the team’s responsibility to ensure that the right ones are used. This is the focus of habit number three; *Research, Test and Pilot Technology*.

To integrate this habit effectively into the BI development approach, we first suggest that the team research a broad range of products and not just what has been heavily advertised in the technical journals. A number of products are available that receive relatively little press exposure - and though some of these products may be risky due to their smaller install base - a smaller vendor offering can be a better choice in practice than a tool made popular through market size and advertising.

“It’s true that there is simply no way to predict every technical problem and performance issue that might arise during a BI project, but that is no excuse for being reactive rather than proactive in dealing with solution limitations.”

We also suggest that the team keep in mind that a collection of tools will be needed. The target of the research should be to develop a short-list of two or three products in each technology area as a ramp-up for a formal test. This includes data movement technologies, reporting tools, portals, dashboards, etc. One shouldn’t assume that a suite of tools from the same vendor is always the best approach. Vendors who offer the best reporting tools may not stack up well on the data movement side and vice versa.

Developing a plan for the formal test phase is important. A formal test consists of scripted scenarios (use cases, test cycles, etc.) under which each product is evaluated and compared. The scenarios must be as real-world as possible; the target hardware, software and the actual users must be involved. The goal is to filter out - under real conditions - those products that don’t fit, don’t work, or are simply inappropriate to the needs of the project.

Once the technology finalists have been selected, the last phase of the assessment is to build a working pilot that brings the pieces together. This does not need to be a large or complex effort, but the pilot must exercise all of the capabilities that will ultimately be brought to bear within the completed system. In our projects the pilot is usually built during the first phase of the development lifecycle. This provides an opportunity to review and assess the viability of the technologies before we proceed too far down the development path. The result of these efforts is a BI technical architecture that is significantly less likely to fail.

Habit #4 – “Acknowledge and Manage Solution Limitations”

We’ve often heard BI practitioners say, “No matter how well you design your DW/BI solution, there will be issues. Maybe the ETL won’t scale as easily as you’d hoped, or the performance of the queries won’t be quite what you expected. Maybe the reporting environment doesn’t offer the flexibility that was claimed. That’s just the way it is.” It’s true that there’s simply no way to predict every technical problem and performance issue that might arise during a BI project, but that is no excuse for being reactive rather than proactive in dealing with solution limitations.

There’s a familiar saying - “If you are not part of solution, you are part of the problem” - and our habit number four, *Acknowledge and Manage Solution Limitations*, reinforces this notion. What does this mean in practice?

First, it tells us that we must recognize that there will be known limits to our BI solution regardless of what tools, methodologies and techniques we use. It forces us to set aside assumptions and complacency and to accept that we will need to develop “what-if” and “what might happen” scenarios to deal with these limits before they occur.

“There’s a familiar saying - “If you are not part of solution, you are part of the problem” - and our habit number four reinforces this notion.”

It directs us to use exploratory techniques such as Capacity Planning, Gap Analysis and Stress Testing to ferret out these limits before they hit the users’ desktops. And it encourages team members to voice the limits of a particular approach early in the development process.

Second, this habit tells us that we must manage these limitations instead of allowing them to manage the team. For example, reporting tools that don’t work must be either made to work or replaced. Users never suffer poor running reports because the team has already put the underlying queries through their paces. And increasing data volumes are not allowed to become a problem because growth has already been anticipated and planned for.

Do we still run into issues of unstable vendor software and poor running queries? Absolutely, but such situations are rare and definitely not normal, and can usually be solved with a simple repair or workaround.

Habit #5 – “Architect for Growth”

During our DW/BI assessment engagements, we have seen a number of client data warehouses that resemble cobbled together science experiments; systems designed without any real growth or maintenance consideration that are difficult to manage and unable to scale beyond initial expectations. We find that many of these data warehouses come about as a result of organizations contracting out small pilot reporting systems that are pressed summarily into corporate-wide service. Technical and information architectures are typically an afterthought with ETL strategies sketchy and dimensional models non-existent. Often these systems stumble along until a technical or data glitch forces the system offline and out of sync with the source data. In time, the prospect of bringing the system back online becomes less worthwhile until it is ultimately written-off.

With 20/20 hindsight, it is easy to dismiss this situation as an example of poor project planning and architecture development. In most cases, such pilot efforts should have been discarded and re-architected rather than pressed into production. But in practice, organizations simply do not have the time or money to “throw one away” and are forced into deploying an immature solution. To avoid these scenarios, our BI team applies habit number five; *Architect for Growth*. This habit advocates that best practices be employed in all of our data warehouse engagements regardless of size and cost. Under this philosophy, no BI implementation is considered an experiment. Since the majority of good DW/BI practices are size and cost agnostic, many can be successfully implemented within projects using inexpensive tools. As such, we find that the fruits of pilot efforts developed with simple design tools (e.g. Microsoft Word/Visio) and inexpensive architectural components need not be discarded as the BI program proceeds.

“Building a data warehouse is as much an exploration as a development process; there are bound to be twists and turns. Even for experienced BI teams it’s often tough to predict what business problems the solution will be asked to solve.”

Does this mean that we advocate SQL, Microsoft Access and Excel as the ETL and reporting tools of choice for a multi-terabyte installation? Of course not, but we do believe that the over-arching best practices employed in building a large, corporate-wide data warehouse are equally valid on smaller, pilot-type scales. As a result, instead of being forced to discard our prototypes and initial efforts, we simply maintain what still applies and upsize whatever has been outgrown. This strategy usually means reduced overall costs, evens out the expense of scaling the initial effort, and significantly reduces re-architecting headaches.

Habit #6 – “Deliver Incrementally and Collaboratively”

Building a data warehouse is as much an exploration as a development process; there are bound to be twists and turns. Even for experienced BI teams it’s often tough to predict what business problems the solution will be asked to solve. We have observed BI teams steadfastly ignore this reality by continuing to cling to obsolete project plans and outdated deliverable lists long after the user community has moved on from the original objectives. In our opinion, this is an outmoded delivery model and a poor way to build a data warehouse.

In recent years we have seen tremendous growth in the adoption of agile to deliver BI projects. One of the things we especially like about agile is that it promotes the idea of adaptive software development. By developing software in an iterative and rapid fashion, a development team can put working software in the hands of the customer quickly to solicit feedback and refine the requirements for the software system under development. To reinforce this approach amongst our team, we promote a sixth habit; *Deliver Incrementally and Collaboratively*.

How does this habit impact our approach to BI development? Here’s an example. One often reads in data warehousing textbooks that the first release of a BI project should focus on a single metric and its associated set of dimensions. Typically this corresponds to a unique cube or star schema, and for many BI projects this becomes the first deliverable. But is this really the best way to introduce the BI solution to the user community?

Why not deliver a single dimension (e.g. “Person”) along with its attributes as a full release? Such a release would undoubtedly prove useful to users long before any data marts or cubes are completed. At the very least, providing the users with access to individual dimensions can stimulate important discussions around error-handling, deduplication and historic data retention much earlier in the design and development process than is typical.

“From the steering committee and project sponsor, to the developers and frontline users, we try to keep everyone involved and excited about the project and committed to making it work.”

To wit, *Deliver Incrementally and Collaboratively* tells our team that we shouldn't think of each release as simply a mechanism for transitioning the system into the users' hands, but rather as an important part of a continuing feedback loop that brings value to the users while supplying direction to the team.

Of course this process must be appropriately managed. For example, we wouldn't suggest a delivery strategy that sees every additional attribute given its own release – but at the same time we wouldn't discount such an approach if it made sense. This illustrates the true value of this habit; encouraging the team's flexibility and adaptability to allow the deliverables to drive the team and not the other way around.

Habit #7 – “Champion the Cause”

BI isn't easy; requirements are often in conflict and project resources are constrained; project plans must adapt quickly to moving targets while technology and integration issues pose scheduling threats; users clamor for more data as batch processing windows become smaller; and at every turn there is the threat of early failure. As pragmatists, we acknowledge the impossibility of managing every possible risk. That's just not realistic. However, we believe that success is more likely if a supportive culture is built into the project. From the steering committee and project sponsor, to the developers and frontline users, we try to keep everyone engaged and committed to the project.

We call the habit that reinforces this belief, *Champion the Cause*. Here are a few ways we encourage this habit within our projects:

- *Designate a lead champion from the business.* This is typically a senior executive who recognizes and believes in the value of the BI vision and acts as the team supporter when times get tough. A good champion commands respect is a terrific project asset.
- *Solicit champions from the user community.* Seek out champions from the user community who can help sell the BI solution. These are users who are particularly excited by the prospects of having access to improved reporting. Accept them into the team by soliciting ongoing feedback and by showing honest respect for their input.
- *Advertise the project's successes.* During the early stages of a BI initiative, it's not uncommon for minor data quality issues to expose significant savings. If the team discovers an opportunity to save the business either time or money, let management and the user community know. This is an immediate return on investment that should be leveraged.

About Quadrus

Quadrus is a recognized leader in IT professional services and solutions. Headquartered in Calgary, Alberta, Quadrus has delivered hundreds of successful projects across Western Canada since 1993. We are committed to providing the highest quality service to our valued clients.

Contact us:

info@quadrus.com
www.quadrus.com
+1 (403) 257 0850

- *Show off the system's extended capabilities.* One exciting aspect of BI is the ability to perform ad-hoc reporting and analytics such as data mining. Showing off the system's potential lets people know what's possible. Champion the capabilities of the data warehouse by performing periodic real-world presentations and seminars that demonstrate the potential power of these techniques. These sessions often prove to be eye-openers that generate a lot of additional interest.

These championing techniques are designed to encourage support of the project through individual efforts and tangible results. In contrast, we sometimes see clients try to build a championing culture through town hall assemblies that extol the virtues of the upcoming solution long before any real results have been delivered. This approach often backfires when the users discover that promised functionality may be months (or years) down the road. As often as not, such well-intentioned cheerleading sessions promote a user-community that is easily disillusioned and difficult to encourage instead of the supportive culture they were designed to build. Championing the solution can only be achieved when the efforts are perceived as sincere and timely.

Good Habits. Great Solutions.

So why do we think these habits are THE habits to use when building BI solutions? One simple reason – They work! Over the many years that we have built BI and data warehouse solutions for clients, our practice has seen technologies come and go. We have observed and participated in industry discussions that have turned longstanding BI practices on their head. We have watched BI grow from a small cottage industry into the multi-billion dollar business it is today. But through all these changes, we've found these habits still hold true; they are just as relevant today as they were a dozen years ago. No matter the client we've serviced, the products we've worked with, or the size of the project we've delivered, every habit continues to drive our teams to the right decisions.

Of course, habits alone are not enough to ensure the success of a project. Success requires equal measures of know-how, experience and good old-fashioned hard work; the traits that every great team must possess. But what our Seven Habits do offer is a leg up on success in the form of guideposts that help point us away from the wrong choices. They've helped keep our BI teams focused on what's important, and we continue to practice each one every day on every project we deliver.