Data appliance platform selection and delivery:
10 pitfalls you must avoid

Authors:
Oliver Halter
Salvatore Passariello
Ritesh Ramesh

pwc

Abstract

Users need to get at more data ever more quickly. Investing in a data appliance helps accelerate data movement across the enterprise. Unfortunately, many enterprises stumble or don't know where to start during their project. The proliferation of data warehouse (DW) appliances brings with it benefits, challenges, and risks.

When it comes to execution, data appliance "replatforming"—often misconstrued as a simple data migration project—can be a complex initiative for many organizations. In this article, we present 10 pitfalls you must keep in mind for such a project.

This article was originally published in the *Business Intelligence Journal*, Vol. 17, No. 2. Reprinted with permission.
Visit tdwi.org for more information.

Introduction

Over the past few years, we have seen the use of DW appliances escalate, and the momentum continues to build. Fueled in part by the growing use of social media and mobile devices, the resulting data explosion is creating a compelling business need for today's organizations—to store and process larger volumes of data and to deliver that data faster, thereby empowering both internal and external stakeholders to meet their analytical and operational needs.

When it comes to data movement, speed is of the essence. In our experience, data latency cascades into a negative impact on business users' ability to access the timely and relevant information they need to make smarter business decisions faster. Organizations can no longer afford to operate in traditional data environments and legacy infrastructures that give rise to performance bottlenecks and stakeholder frustrations. In an era where enterprises are focused on figuring out creative ways to help operational and strategic decision makers leverage business intelligence, management is increasing its emphasis on achieving speed of data movement across the enterprise.

Most DW appliances currently available offer an integrated stack of hardware, software, and storage in a "single box," with a scalable parallel-processing framework capable of faster data processing and information delivery. These offerings are designed to provide customers with distinct cost, processing, and operational advantages over and above the advantages delivered via the traditional model, where core components are individually procured, configured, and maintained.

The decision to invest in a data appliance as a business enabler for accelerating data movement across the enterprise is often a key, positive step. That being said, IT organizations struggle with complex delivery and execution challenges that arise during the integration of their existing technology stack—including data integration tools, reporting tools, custom applications, security processes, and operational standards—with the new data appliance. Poor planning and execution can lead to years of extended wait times before organizations can realize the benefits of the new data appliance.

To guide your navigation around these challenges, we present 10 potential pitfalls you must avoid when you select a data appliance platform. We offer suggestions for each to help you mitigate risks, overcome problems, and meet your short- and long-term project objectives.

Finding the "best-fit" solution requires the investment of sufficient time and effort up front, involving internal stakeholders in selecting vendors before issuing your RFP.

Pitfall #1: Initiating the vendor-request process without first involving internal technology and business stakeholders

Data appliances are not silver bullets for solving all the business problems you may encounter. Because every organization is different, it is essential to evaluate appliances to determine whether they are a good fit for the specific business and technology needs of your unique enterprise. Finding the "best-fit" solution requires the investment of sufficient time and effort up front, involving internal stakeholders in the vendor-selection process before issuing your request for proposal.

This collaborative approach brings a wide range of issues, expectations, and priorities to the surface, which must be discussed and documented—including the specific applications, tools, and business-unit requirements that the new appliance must meet. A holistic vendor-selection process typically results in a common, aligned vision that, in turn, enables the organization to manage the expectations of its internal stakeholder base right from the start of the data appliance replatforming strategy.

Pitfall #2: Assuming that the proof of concept is all about executing fast queries and data loads

One critical success factor for the vendor pilot is the capability of the new appliance to seamlessly integrate with the tools and applications that already exist in the organization. Another factor is its ability to scale and load large volumes of data to execute user queries faster. Doing your homework to gain a real-world perspective is essential. Many replatforming efforts fail during the execution phase because organizations often do not realize until very late in the project how many years it will take to effectively integrate their current applications to the point that they work seamlessly with the new appliance. Failure to successfully integrate the new appliance with the organization's existing tools and applications limits your ability to achieve the expected return on investment (ROI) and realize the strategic business benefits in a shorter period of time. It also often makes it impossible to replace legacy DWs.

Once you have pared down the number of multiple vendors in initial rounds of the vendor-selection process to a short list of two, it is time

to extend the evaluation scope of the pilot to simulate a thin, horizontal slice of the current operational environment in the new appliance.

Replatforming is not just about migrating data from your old database platform to the new data appliance.

We recommend that you create benchmark scenarios to evaluate the capability of the appliance to (1) seamlessly and simultaneously integrate across various tools and applications in the current technology stack, (2) load high volumes of data, and (3) execute queries more rapidly. Make sure that the scope of this pilot phase spans such vital areas as key points of integration in the current environment, workload management, security, backup and restore capabilities, and user inputs.

By investing additional time and resources in the pilot phase as an extension of the proof of concept, organizations can detect any critical functional/nonfunctional requirements, processes, and organizational gaps inherent in the appliance being considered—a vital step, given that replatforming efforts often amount to a multi-million dollar investment.

Pitfall #3: Viewing the replatforming effort as a tactical IT project rather than as a long-term transformation effort

Many large organizations find DW appliance replatforming to be a complex initiative. By their very nature, replatforming efforts affect several existing critical touch points within the organization—people, processes, and technology—and require change and transformation as part of the initiative. Intricate interdependencies exist among these core components that will eventually drive the costs, planning, and execution sequence of the replatforming effort.

Recognize that replatforming is not just about migrating data from your old database platform to the new data appliance. Its scope is much more complex. To work effectively with the new appliance, a redesign of existing data integration, reporting, and custom tools and applications—and underlying technical architecture—is likely to be required. Organizations should view any replatforming program as a strategic IT transformation initiative—delivering incremental benefits for the enterprise with specified milestones rather than one tactical "big bang" IT project—and plan accordingly.

Pitfall #4: Underestimating the required investments in project management, governance, and change management

Many organizations initially view DW appliance replatforming as a simple IT initiative to enhance their existing technology portfolio with new parallel and scalable database technology—and they expect to accomplish this while minimally affecting their current technical environment.

However, a replatforming effort differs from a typical data warehousing initiative in size, scope, and complexity. It always involves decommissioning resources in the existing technical environment while simultaneously developing new functionality. A typical replatforming environment involves a heterogeneous mix of people, processes, and technology across the organization—for which the scope, priorities, and interdependencies must be successfully managed to meet program milestones. As changes occur in the existing technical environment, effective management across the organization is critical to success.

All too often, organizations that underestimate the scale and complexity involved in a replatforming initiative hesitate to invest sufficient resources in the planning phase, project management office, and governance functions—which leads to program failure during the execution phase. The establishment of a dedicated, full-time project management office with an effective governance and change-management function is a critical success factor for any DW replatforming effort.

Pitfall #5: Building a short-term, tactical organizational structure incapable of sustaining scale and complexity over time

A replatforming effort entails multiple vertical and horizontal work streams comprising technical and functional subject matter experts, technology specialists, external vendor resources, and work stream leads. We recommend that you organize and structure each work stream with technology and subject matter experts that are aligned, either horizontally or vertically, with the milestones of the program. In the early stages, it is critical to involve specialized resources with expertise in the new data appliance and then leverage their depth of knowledge and experience about standards setting, leading practices, execution guidance, and knowledge transfer.

Organizing noncollaborative, ultitasking teams with short-term tactical goals is not an efficient and sustainable approach in the long run. As scale, complexity, and change build over time, such teams typically lose focus and fail to accomplish their execution tasks. Clearly, then, the iterative process of building, reconfiguring, and managing teams aligned to program milestones is critical and therefore should be designated as an integral, centralized function of the project management office for any replatforming effort.

Pitfall #6: Underestimating the need for architecture and design leading practices and oversight

Leading practices, design standards, and the complexity of executing a data appliance replatforming strategy can vary based on several key factors, depending on whether:

- The new platform being selected comes from the same technology vendor as the existing platform or from a different technology vendor
- The new platform follows a row-based or a columnbased data partitioning scheme
- The new solution uses a simple data migration or a complex data conversion strategy and re-architecture

Because bad decisions made during the architecture and design phases can cause the overall replatforming program to fail, you must establish a centralized architecture board and design a review organization comprising technical experts. These experts can be invaluable in establishing standards and leading practices, as well as in the review of and sign-off on all key design and architectural decisions made across all work streams involved in the project.

Ideally, an organization will hire an architect with knowledge and experience in the new technology who—in tandem with the existing architects with organizational and historical knowledge—can effectively guide critical decisions. This collaborative approach will help to ensure that the teams involved in the process adhere to a common, consistent framework and to leading practices for the new data appliance during execution, thereby contributing to the program's overall success.

Pitfall #7: Implementing the solution in a "big bang" waterfall approach

It is a common misconception that because replatforming programs are often large and complex, agile execution and ideas do not apply. Although the planning, governance, and change management functions do need to be organized through the project management office (which acts as a central function for the program), once work streams with clear, measurable goals and milestones have been established, individual work streams may use agile methodologies and leverage leading practices to expedite delivery functions. Individual team leaders should encourage their respective teams to think "big picture," collaborate, start small, iterate, and then evolve toward their final solutions and deliverables.

During the design phase, it is especially important to identify areas of automation and reuse wherever applicable so that the team can effectively manage resource-intensive areas that require application code to be redesigned and rewritten. Currently available third-party tools can facilitate the automation of routine tasks and functions that are required to move data and code from the old database platform to the new appliance platform. We recommend exploring these tools with an eye toward putting them to work during the execution phase.

These days, teams involved in a replatforming program operate in a dynamic technical environment where change is constant and rapid. It is critical that they approach the design process with change, reusability, and the iterative review and release of work products top-of-mind. By adopting key principles from the agile framework for execution and delivery, savvy delivery teams can contribute to the overall success of the replatforming program.

Pitfall #8: Underestimating the training investments and cultural transformation required within the organization

An organization that has embarked on a journey to invest in a DW appliance should be able to clearly articulate the vision, goals, and business drivers of the replatforming effort. Involving teams early in the process—clearly communicating the program's goals and objectives and what they will mean to employees and other stakeholders—will help teams understand the big picture, boost morale, and generate active support of and participation in the effort.

A certain degree of cultural transformation is required of the teams as part of adapting the life cycle to the new data appliance platform. It is vital for program leaders to engage and communicate with teams—not only up front but on a regular basis—clarifying the program goals, reporting progress against milestones, and providing updates about priorities, wins, and the status as things change.

We have seen organizations that fail in this regard suffer resistance and higher attrition rates among their

employees who, not having been brought into the loop, view the effort as a pure technology change that the organization has forced upon them. Further, substantial investment in the areas of training and knowledge management within the organization is vital, as this investment enables teams to learn, share, and apply their new skills on the job, ultimately resulting in improved productivity across the company.

Pitfall #9: Designing a program structure without building in key performance indicators (KPIs) to measure, monitor, and reward success

Measuring and monitoring success in an appliancereplatforming effort starts with establishing measurable milestones in the program. It is important to align and map each milestone to a specific set of business drivers and benefits—either quantitative or qualitative—that will be delivered at that milestone.

- ***Quantitative benefits*** are easily measured. Examples include improvements in application run time and/or data load times, decreases in data latency in the processing life cycle, and increases in query performance. Quantitative benefits directly contribute to the "speed of data movement" from source to target, enabling the enterprise to make smarter business decisions faster. It is critical to translate technical and process improvements (such as speed of queries, loads, etc.) into tangible business benefits that can be derived from improved processes.
- ***Qualitative benefits*** must also be tracked and measured, although this is more difficult due to their more intangible nature. Examples include worker productivity and satisfaction, better end-user experience with reports and dashboards, and improved technical system policies and procedures.

The project management office must measure, track, and capture both quantitative and qualitative performance metrics as a means

of rewarding desired behavior and retaining valued employees. In addition, many enterprises use these metrics as a source of input for their performance scorecard system to tweak program milestones and continually improve the program. Over time, performance metrics and milestones must be periodically reviewed and communicated to the teams and stakeholders involved in the program.

Pitfall #10: Overlooking hidden costs

Apart from the benefits of faster data processing and information delivery, most DW appliance vendors project substantial cost savings on infrastructure due to their inherent ability to offer an integrated stack of hardware, software, and storage components in a "single box" for a fixed dollar cost. Furthermore, because fewer personnel are required for the administration and management of the appliance platform, ongoing maintenance costs are projected to be substantially lower.

That said, there are many overlooked service-related costs after procurement, such as training and knowledge management costs, external vendor-services costs for executing the replatforming effort, additional future costs for storage and processing components, hardware/software decommissioning costs, and personnel costs for engaging internal subject matter experts from other business units in the program. To reduce the likelihood that any surprises will arise to adversely impact the execution life cycle of the replatforming effort, leading organizations should proactively analyze and account for these types of hidden costs in their program budget and cost-benefits plans.

Moving ahead

When it comes to execution, data appliance replatforming— often misconstrued as a simple data migration project—can be a complex initiative for many organizations. Replatforming affects several existing critical touch points across the organization's people, processes, and technology. In addition, it requires change and transformation during the execution of the initiative.

Before investing in a DW appliance as a key business enabler (in terms of accelerating data movement across the enterprise while reducing expenses), it pays to invest significant time and effort to complete sufficient due diligence. Leading organizations view DW appliance replatforming as an IT transformation program designed to deliver incremental benefits to the enterprise with its milestones, and they plan their approach to execution accordingly.

Focusing on, and adequately investing in, these 10 risk areas will position your organization to successfully execute on its IT strategy, accelerate data movement across the enterprise, attain the expected ROI, and, ultimately, fully reap the benefits of the new DW appliance.

The authors

Oliver Halter, Partner, PwC Advisory Information Advantage practice, oliver.halter@us.pwc.com. Oliver specializes in Enterprise Information Strategy and Architecture and in creating tangible business value from information.

Salvatore Passariello, Director, PwC Advisory Information Management practice, Salvatore. passariello@us.pwc.com. Sal focuses on Enterprise Data Management, Information Architecture, Enterprise Architecture and Business Intelligence.

Ritesh Ramesh, Manager, PwC Advisory Information Management practice, ritesh.ramesh@us.pwc.com. Ritesh specializes in high performance data processing technologies, data warehouse appliances and scalable database architectures.