PART ONE

Planning Your Customer MDM Initiative

CHAPTER ONE

1 Defining Your MDM Scope and Approach*

Success is not final, failure is not fatal: It is the courage to continue that counts.

—*Winston Churchill*

MDM APPROACHES AND ARCHITECTURES

Master Data Management (MDM) is about bringing master data together to facilitate the employment of master data management services—such as data governance and stewardship; data quality, metadata, hierarchy, and overall data lifecycle management—and ultimately, to serve as the single source of truth for the business. Customer MDM focuses on the *customer* data domain in particular and its associated properties, such as company name, tax ID, addresses, contacts, accounts, and company hierarchy.

In addition to data domains, such as customers, products, partners, and suppliers, data inside a company can also be classified as operational or

* Portions of this chapter are based on Dalton Cervo's contribution (Chapter 12: Master Data Management) to Phil Simon's book, *The Next Wave of Technologies: Opportunity in Chaos,* Hoboken: John Wiley and Sons, Inc., 2010.

nonoperational. Operational data is the real-time collection of data in support of a company's needs in their daily activities. Nonoperational data is normally captured in a data warehouse on a less frequent basis and used for business intelligence (BI). This particular classification of data is relevant in this context because it can be used to distinguish most common MDM initiatives.

Although the very essence of implementing MDM is in the appliance and fine tuning of MDM practices to fit the enterprise architecture and business model, MDM implementations as a whole can generally be categorized into three major types of initiatives based on its primary focus being operational or nonoperational data:

1. Analytical MDM: address BI
2. Operational MDM: address business operations
3. Enterprise MDM: address both BI and operations

Each has a somewhat different objective and carries distinct levels of complexity, risk, and impact. Companies should perform detailed analysis to decide which approach is required. At a minimum, an MDM program must take into consideration business and IT requirements, time frame, resource availability, priority, and the size of the problem to be addressed.

Deciding which approach to implement is dependent on the business case, which is explained in more detail later in this chapter. Because each of the previous approaches targets a different category of information, they ultimately impact a company at varying degrees. Figure 1.1 depicts the level of intrusiveness of each MDM approach.

Operational data is inherently more critical to a company than nonoperational data due to its usability and timeliness. Therefore, analytical MDM is the least intrusive approach, followed by operational MDM and obviously the all-encompassing enterprise MDM, which is a combination of both analytical and operational MDM.

Naturally, more intrusive MDM projects involve both higher risks and higher likelihoods of disrupting companies' daily operations. It is important to notice that the figure does not suggest a sequence or phases to be adopted when implementing an MDM solution. As a matter of fact, phased deployments need to be observed from two different perspectives. One is concerned with progressing from one approach into another, such as starting with an operational MDM, then an analytical one to complete the enterprise solution. Another way to look at phased deployments is within a particular approach. It is not uncommon to start an operational MDM integrating just a few legacy systems,

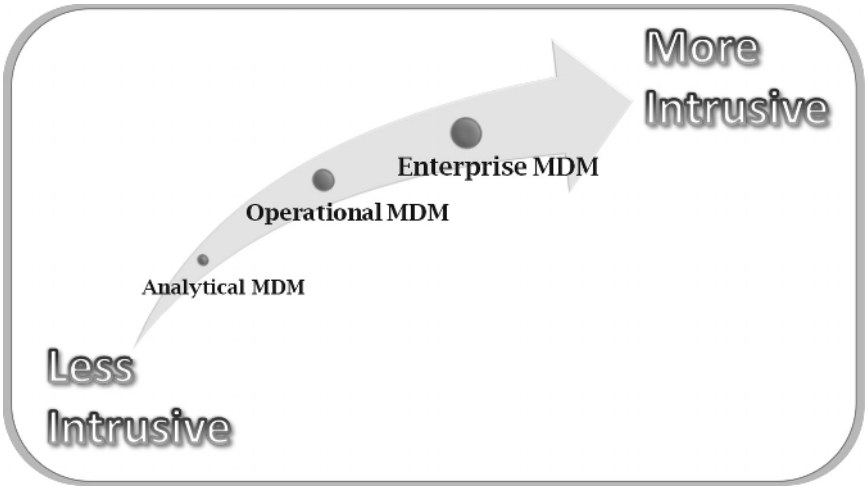

FIGURE 1.1 MDM Approaches

and slowly incorporate others. More about phased deployments will be discussed in Chapter 3.

Next, each of the approaches is explored further with the most common architectures employed for each of them. Keep in mind these are generic frameworks. MDM can be so encompassing and pervasive that the number of potential combinations can be many. Hybrid solutions are also very common. Finally, many subjects in the MDM arena don't have a universal terminology. What is called approaches and architectures in this book may be called styles, framework, or implementation in other books along with other varying definitions. What is important is to understand how the master data is integrated, used, maintained, improved, and governed.

Analytical MDM

Historically, analytical MDM has been the most commonly adopted MDM approach. This stems mostly from the relative simplicity of leveraging data warehouse projects. It is beyond the scope of this book to describe data warehouses in detail, but the following summary of the three primary data warehouse architectures might help you understand how MDM projects can benefit from this already existing integration:

1. **Top-down.** Major proponent is Bill Inmon. Primarily characterized by a data warehouse as a centralized and normalized repository for the entire

enterprise, with dimensional data marts containing data needed for specific business processes. Up-front costs are normally higher and it takes longer initially until common structure is designed, built, and sources are integrated, but it is more adaptable afterward.

2. **Bottom-up.** Major proponent is Ralph Kimball. Data marts are first created to provide reporting and analytical capabilities for specific business processes, and can eventually be integrated to create a comprehensive data warehouse. Provides results quickly to each independent business unit, but overall data integration is potentially harder to achieve.
3. **Hybrid.** A combination of top-down and bottom-up approaches, characterized by a high-level normalized enterprise model, more quickly integrated with business specific data marts for faster results.

One may ask: If there is already a data warehouse integrating the data from across the enterprise, isn't that MDM? The answer is: not necessarily. It actually depends what is being done with that data. Bringing the data together is just one piece of MDM. The other piece is applying MDM practices, such as identity resolution; data cleansing, standardization, clustering, consolidation, enrichment, categorization, synchronization, and lineage; metadata management; governance; and data stewardship.

Bottom line is a data warehouse, and data mart infrastructure can work as the conduit to a much larger and encompassing MDM program. Conversely, the business intelligence, analytics, reports, and other outputs relying on the data warehouse and data marts will greatly benefit from the additional practices imposed by MDM—above all, data quality and hierarchy management improvements. Keep in mind that in this context, a strategic or a tactical BI implementation is implied instead of an operational BI since the underlying data is nonoperational.

Figure 1.2 depicts a common architecture adopted by companies implementing an analytical MDM approach.

Figure 1.2 shows that an extract, transform, load (ETL) process gathers data from disparate operational systems. Ultimately, the data is stored on an enterprise data warehouse (EDW). EDW and associated data marts become the source of master data for BI and analytics. Since EDW is now a single source from an analytical perspective, it is also the centerpiece for what can be called MDM services.

Analytical MDM is the quick-hit approach. While companies can quickly make a tremendous impact with respect to reporting and BI, with the analytical MDM approach relatively minimal inputs yield corresponding outputs.

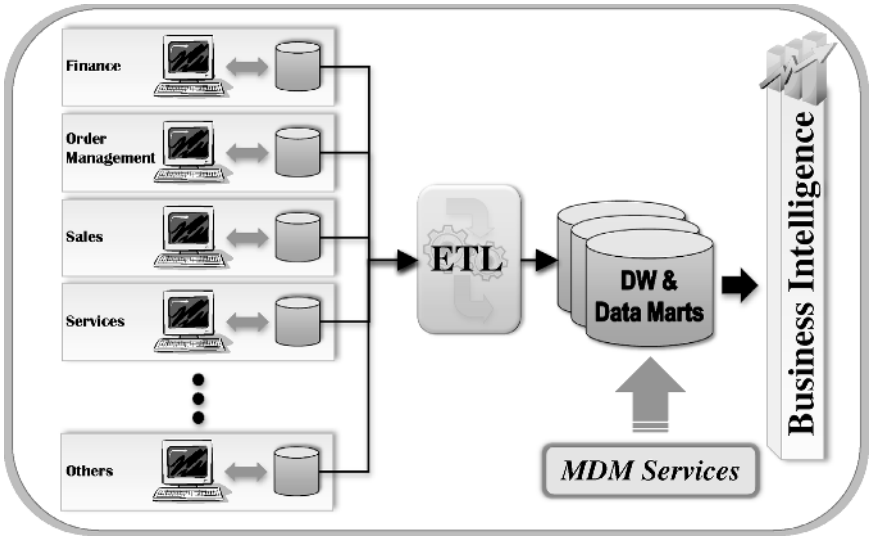

FIGURE 1.2 Analytical MDM

Specifically, companies fail to harvest the benefits of the MDM services back to their operational data. Remember, the data improvements are happening in the data warehouse, which is downstream from the operational systems. What's more, the analytical MDM approach does not enforce any regulatory or audit requirements since those are mandatory at the operational level.

Another drawback with this implementation is the possibility of adding one more fragmented and incomplete data system to the company. Obviously, the quality of the results will be directly related to the quality of the MDM services applied to the data. But a less obvious conclusion is the quality of the results is also directly related to the amount of data sources integrated. Certain lines of business (LOBs) are very sensitive about feeding data warehouses with their operational and strategic information, making it hard to achieve comprehensive integration.

On the other hand, it is possible for companies implementing an analytical MDM to influence the operational world. Analytical teams have access to an integrated view of the data and its underlying quality. They can recognize bad data and potential root-cause offending practices relatively quickly, as well as correlate discrepancies across LOBs. This is powerful knowledge that can be used by a strong data governance team to influence and improve data quality and business practices at the source. Be aware, however, that operational LOBs

tend to be very resistant to this approach and to succeed with this practice, strong sponsorship from high-level executives is necessary.

Operational MDM

Operational MDM targets operational systems and data. It provides the opportunity to consolidate many, and ideally all, disparate operational data systems across the company, and become a true system of reference. This is obviously an enormous task. From a data integration perspective, the difficulty increases with the volume of data to be integrated along with the level of disparity among the systems to be combined. But it is much more than simply data integration. It is about business process integration and massive technological infrastructure change, which can impact virtually everyone in the company.

Depending on the size of the company, an operational MDM will likely be deployed in phases. Breaking down what is included in each phase can vary widely as well. One method for phased deployment is gradually migrating each data system into a single MDM repository until all systems in scope have reached end-of-life (EOL).

Another method for breaking down phases is gradually migrating portions of data from a single system. Sometimes this is necessary because it is not possible to promptly EOL a particular legacy system if not all its business processes have been transitioned to the new application yet. It may sound strange that there is a need to start transferring the data if the system is still operating. But that is sometimes necessary to support other already migrated systems that have dependencies on that particular legacy data.

Finally, a combination of both phased methods are not uncommon either, with systems and portions of data making their way to the single MDM source at contrasting techniques. The data integration component of MDM is obviously complex, and companies need to be very creative in finding the best method for consolidating legacy data.

The bottom line is that it can be very difficult to EOL a given operational system because it is not only a technical issue; it is a business issue, as well. Changing business practices that have been in place for years and years can be overwhelming. Besides, it could impact customer relations, and that is the last thing anyone would like to happen. Therefore, to avoid disruption of current business practices, a common practice is to implement a temporary interface between the legacy and the new system until the transition is finally complete.

Chapter 3 will get into more detail regarding phased deployments, data migration, business process reengineering, build versus buy MDM, and so on.

Nonetheless, once an operational MDM is implemented, companies can leverage it into the analytical world for a complete enterprise MDM solution with relative ease. Operational MDM can be accomplished via three different architectures:

1. Single central repository architecture
2. Central hub and spoke architecture
3. Virtual integration

Note that a service-oriented architecture (SOA) with an enterprise service bus (ESB) and business process orchestration is not required to make the MDM repository or the federation system available, but it is the most common and effective architecture.

Single Central Repository Architecture (SCRA)

In this architecture, a single central repository within the operational environment serves as the source of data to an integrated suite of applications and processes. Only one physical copy of master data exists.

It is important to emphasize that this approach may obviate the need for certain applications. In other words, after the consolidation of data, a company may not need all of its previous applications. Required applications dependent on that data might need to be rewritten, or will likely require some interface or other major changes to maintain integration.

SCRA guarantees consistency of master data. However, it can be very expensive—if not impossible—to implement due to potentially inflexible off-the-shelf applications in use (although, if reached, this could actually be the easiest and cheapest to maintain). SCRA could potentially require a massive data conversion effort, depending on the size of the company and the number of disparate systems.

In Figure 1.3, multiple legacy systems go through a data conversion step to bring data into a central hub. This conversion normally takes place in phases to minimize impact and lower risk of concurrently converting multiple legacy systems. When the central hub is operational, it is then used by application systems that would either replace legacy systems or add new functionality to the company. In this particular case, new application systems do not have their own versions of master data.

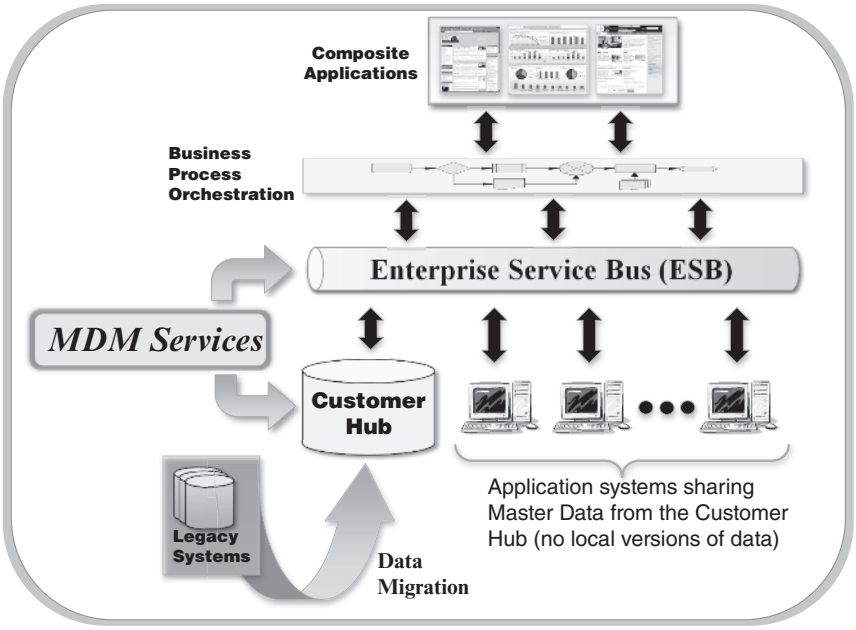

FIGURE 1.3 Single Central Repository Architecture

Central Hub and Spoke Architecture (CHSA)

This is a more common variation of SCRA. Like SCRA, CHSA has an independently deployed common repository. However, CHSA does not require that all applications and processes are fully coupled to the hub.

The major advantage of this architecture is the efficiency of a central hub hosting the master data, combined with the flexibility to support spoke systems operating relatively decoupled. This flexibility is important when integrating commercial, off-the-shelf (COTS) applications with an MDM solution.

Some of the applications can act as spoke systems with independent data models, but cross-referenced and synchronized to the central data. To be sure, CHSA alleviates some of the problems presented by SCRA, but CHSA can still require a massive data conversion effort and new interfaces between the hub and its spokes.

In Figure 1.4, multiple legacy systems go through a data conversion step to bring data into a central hub. Again, this conversion normally takes place in phases to minimize impact and lower risk of concurrently converting multiple legacy systems. When the central hub is operational, application systems then

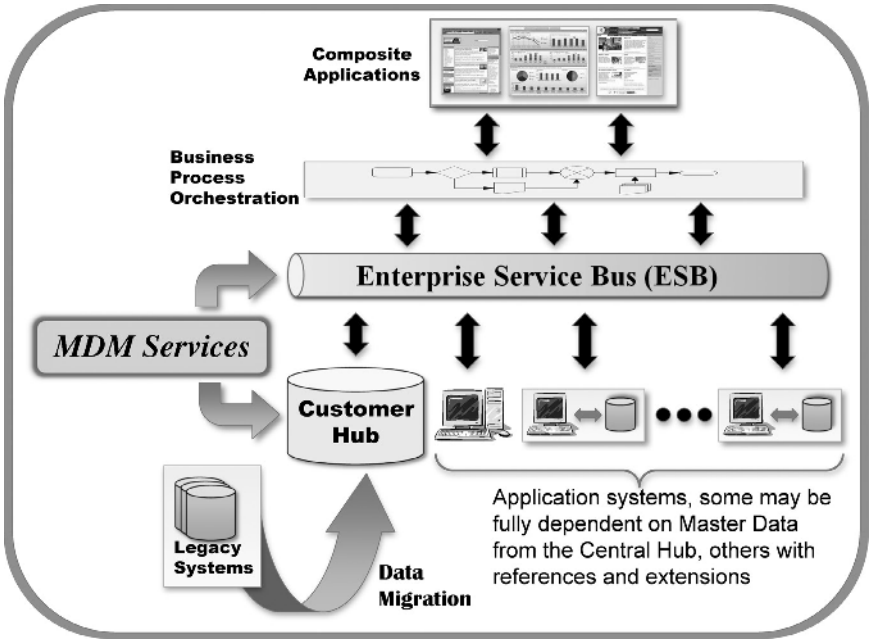

FIGURE 1.4 Central Hub and Spoke Architecture (CHSA)

access it to either replace legacy systems or add new functionality to the company. Spoke systems are synchronized and integrated with the central hub.

Virtual Integration (VI)

Virtual integration is a generic term to represent solutions that don't physically copy existing data into a new repository. This is a fundamental difference compared to the previous two methods. Registry and data federation (DF) are common VI architectures. A VI system aggregates data from multiple sources into a single view by maintaining a metadata definition of all sources. Data across multiple sources is collected in real time through some pre-established keys connecting the VI system and its sources. MDM services are applied to the dynamically collected data, becoming a new source of trusted data to downstream process applications. DF systems normally provide a more robust infrastructure than a simple registry implementation.

The biggest drawback with this implementation is the lack of data improvement propagation back to the source. VI provides benefits to consumers of its services, but not to the original sources of the data. Conversely, due to

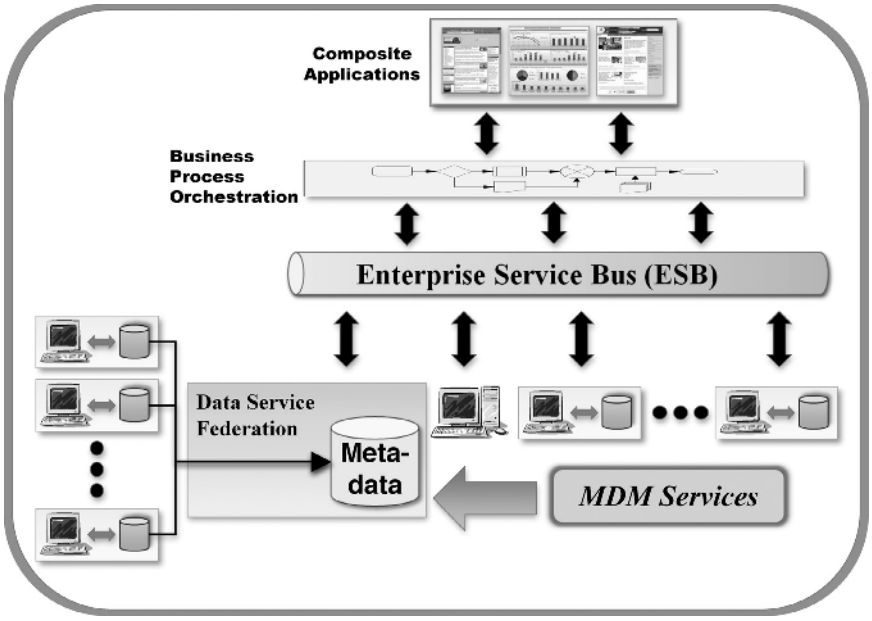

FIGURE 1.5 Data Service Federation (DSF) System

its nondisruptive nature, it is relatively simple to deploy. It could be a good first step before embarking into a central hub implementation.

In Figure 1.5, a data service federation (DSF) system collects real-time data from multiple existing sources. The data is not physically copied into the federated system. Information about the data in each source is stored on a metadata repository. It is further used to determine which system and data element to access based on requests performed to the DSF.

Enterprise MDM

Enterprise MDM is a combination of both operational and analytical MDMs. As such, it can be implemented by combining the architectures previously discussed.

A data warehouse solution could be added to any of the three operational MDM architectures. As an added bonus, most of the MDM services that would be needed in the warehouse are already functional in the operational system, making the maintenance of your data warehouse much easier. Furthermore, the ETL function of the analytical MDM should be much simpler since companies now maintain fewer systems from which to extract data. What's more, the data should be cleaner, standardized, and already consolidated.

Data federation offers another potential solution. DF could be expanded to provide a view into multiple departmental data warehouses in addition to operational systems. Through this method, DF becomes the single point to resolve complex BI queries. This solution reduces both companies' costs and complexity by lowering the need for an extra and expensive database server. However, there's no free lunch here.

DF technology takes a toll on performance of the operational and transactional data sources that it queries. It requires that transactional data sources are always on. This is in stark contrast to batch load data at preset and convenient times as normally done by data warehouse implementations—for example, at 4 A.M., while few users are accessing the system. BI queries can be quite complex and aggregate a multitude of data. Data warehouses are normally optimized to support those queries, making a DF implementation for this purpose potentially unfeasible. If companies go this route, then they should proceed with caution and perform extensive load testing to confirm viability.

Figure 1.6 shows one possible enterprise MDM architecture implementation.

In conclusion, the number of combinations of MDM approaches and architectures is large. The previous figures and categories are meant to be

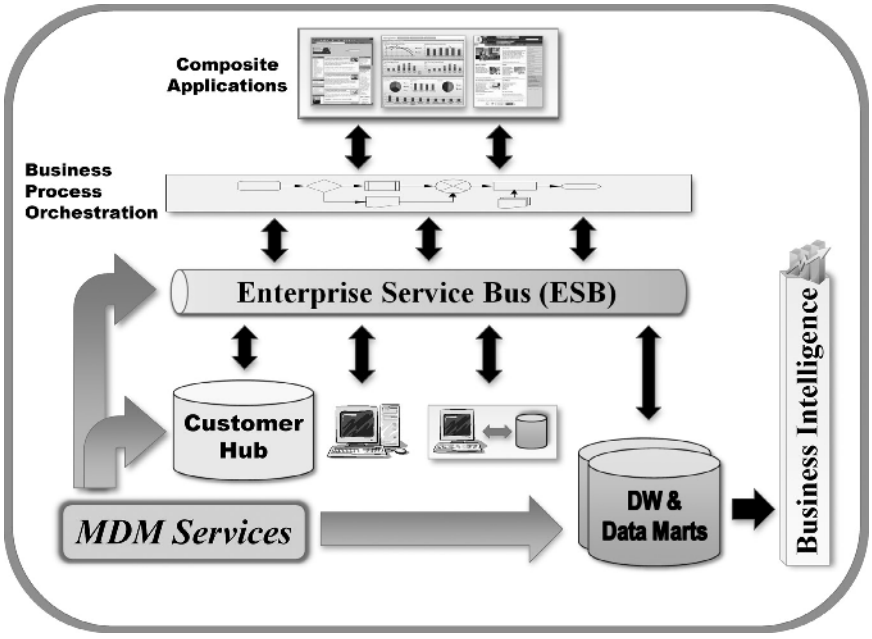

FIGURE 1.6 Enterprise MDM

general guidelines and the most common implementations. It is important to consider the data domain in scope (obviously *customer* in the context of this book), the purpose of managing the data (operational or analytical), and the technical architecture (central hub, data warehouse, virtual integration, hybrid).

Next is a description of the common business cases normally utilized to justify the deployment of a Customer MDM solution, followed by which approach(es) and architecture(s) best fit each of the business cases. Concluding this chapter is a discussion around the elusive ROI question.

DEFINING THE BUSINESS CASE

Why is MDM important? Quite simply, MDM not only gives companies the opportunity to better manage their key data assets and thereby improve the overall value and utility the data provides internally, but it also exposes internal process issues and business practices (or lack thereof) that are the underlying constraints to having and maintaining good data. Often, these underlying issues are generally known, but it's not until an MDM initiative is launched that the various business teams can or will start effectively addressing the issues. Looking at this from a Customer MDM point of view, the lack of having well-orchestrated data management practices will typically result in one or more of the following risks:

- Increased costs due to operational and data redundancies or differences across lines of business.
- Higher risk of audits and regulatory violations.
- Poorer BI and analytics, adding to customer frustration and missed opportunities.
- Customer/partner/vendor/employee dissatisfaction and consequently unrealized revenues.
- Possible overpayment of vendors and customers stemming from duplicate records.
- Over or under delivery of customer services due to inconsistent customer identity and tracking.

MDM does much more than just bring data together. It involves an entire set of processes, services, and policies that go along with it. Most MDM experts agree that the three main reasons used to justify an MDM implementation are cost reduction, risk management, and revenue growth, as shown in Figure 1.7.

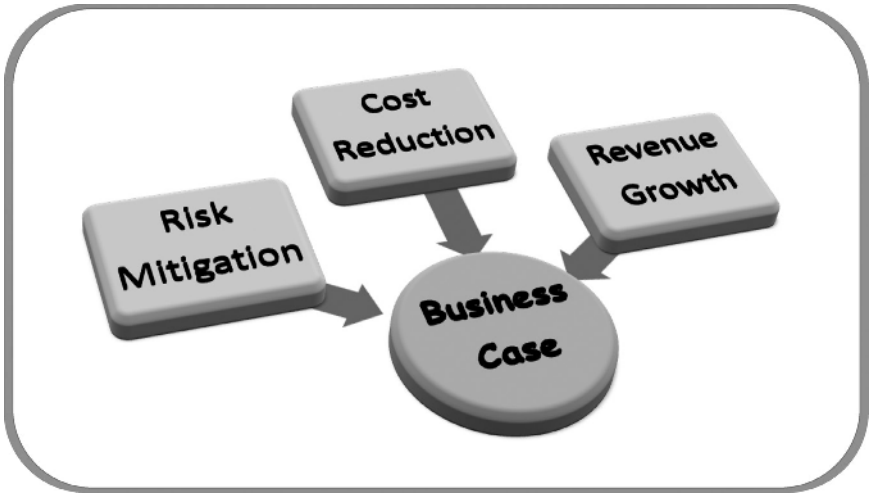

FIGURE 1.7 Business Case for MDM

Cost Reduction

From an information and knowledge perspective, customer data is the nervous system of a company. It is extremely pervasive and impacts virtually every LOB in the company. It is key information associated with transactions from order to shipment, and associated services before, during, and after a deal. Furthermore, proper management of information requires a complex infrastructure, trained personnel, and a set of well-defined and properly followed business processes.

There are costs associated with hosting and manipulating data, as well as costs associated with the activities relying on the accuracy of the information itself. Therefore, any effort targeted at any of those elements will help lower cost. Data hosting is about IT resources, data manipulation is about business processes, and accuracy of information is about data quality. MDM is a perfect fit because it addresses all these facets.

From a data hosting perspective, MDM emphasizes data consolidation through the elimination of redundant systems. This type of activity lowers IT capital expenditures, reduces costs associated with software licenses and maintenance fees, decreases spending on human resources due to having fewer systems to operate and maintain, and eliminates certain consulting fees as duplicated systems are retired.

From a data manipulation perspective, MDM stresses the need for more efficient and effective business processes. A strong data governance program

facilitates the standardization of business rules, policies, and procedures, therefore increasing the predictability of data changes needed. Data stewardship becomes better equipped to carry on data operations, lowering operating costs related to inefficient business processes, and minimizing the number of workarounds and redundant tasks.

Finally, data quality is a core competency of MDM. Data quality practices focus on cleansing, standardization, consolidation, and enrichment of information, therefore lowering costs related to incorrect business decision based on incomplete or wrong information. Better information also helps lower costs related to delivery errors, shipping fines, inventory management, wasted direct-marketing catalogs and other marketing initiatives, and so on.

Risk Management

Companies are susceptible to all kinds of risks. Large portions of them are dependent on having precise and timely information. For example, incorrect inventory management can lead to loss of customers, while improper compliance management can lead to lawsuits and hefty fines. Strategic objectives based on poor data quality analysis and results can have drastic effects, from missed targets to a total collapse of the company.

When looking at one aspect of risk, companies have legal and ethical obligations toward society. Law and ethics are intertwined, and a commitment to both is necessary to guarantee proper operations and successful business performance. Properly tracking and monitoring compliance to these rules requires active information management.

Obviously, the amount of regulation varies by industry, but just about every company will be subject to a multitude of labor laws and compensation, privacy at multiple levels, contract management, pricing, and so on. Additionally, regulatory compliances such as the Sarbanes-Oxley Act (SOX), U.S. Patriot Act, Basel II, and other financial related rules must be followed and flagged immediately if violated. These are all data-driven violations that threaten the survivorship of a company.

Looking at another aspect of risk, proper business decisions based on correct information will certainly lead to more opportunities and less waste related to mistakenly assigned and/or improperly executed activities.

Anyhow, the proper management of data as an asset will increase the probability of positive risks, such as better strategic decisions and increased opportunities, while minimizing the likelihood of negative risks, such as frauds, lawsuits, audit findings, loss of certifications, fines, penalties, and so on.

Revenue Growth

Companies are constantly looking for opportunities to increase revenue. While there are many strategies and variables to make it happen, it all usually starts with proper knowledge and understanding of the driving forces affecting that particular industry. These driving forces include: customer needs, market maturity, competition, partner relationship, product, price, distribution, promotion, and so forth.

A very common and effective strategy is improving customer relationships. However, this is easier said than done. Companies have been implementing Customer Relationship Management (CRM) systems to achieve that goal, but results are not necessarily optimal. A CRM system is just as good as the data behind it. If the data is siloed, and does not encompass the many LOBs across the company, it is difficult to have a clear understanding of the customer. This gap prevents a company from having a full insight into the customer information, or what is also referred to as a 360° view of the customer. This incomplete information is a double-edged sword, preventing a company from fully understanding its customer needs as well as preventing itself from being more efficient and effective when planning the next move. MDM and its practices address the root cause of those issues, consequently enhancing customer satisfaction, lowering customer churn, increasing sales, and ultimately improving revenues.

Better information will also lead to better marketing campaigns, improved partner relationship, and supply channel management.

Finally, an improved understanding of market forces, customer needs, and a company's own strengths and weaknesses will lead to enhanced strategic decisions, eventually culminating with a better focused company and potentially tactical mergers and acquisitions.

SELECTING THE RIGHT MDM APPROACH

As discussed previously, companies typically use three main reasons to justify an MDM implementation: (1) cost reduction, (2) risk management, and (3) revenue growth.

But once a company selects its main area of focus for its business case, how does it decide which MDM approach to implement?

Enterprise MDM represents a combination of analytical and operational MDMs. To be sure, companies can use enterprise MDM to solve virtually any data and integration problem. Furthermore, if a company does want to address

TABLE 1.1 Business Rationale for Different MDM Approaches

Business Case	Recommended MDM Approach	Rationale
Risk mitigation	Operational MDM	The biggest concern is control. Implementing MDM on a data warehouse (analytical MDM) will not help, because operational data must be regulated to minimize risk and increase compliance.
Cost reduction	Operational MDM (maybe Enterprise MDM)	Most likely, the majority of costs are related to operations (see section on the business case for MDM). This means that an operational MDM is sufficient. However, inconsistent and incorrect data may also have a huge cost impact on wasted marketing campaigns. As a result, enterprise MDM may be required, depending on the situation.
Revenue growth	Analytical MDM	Revenue growth is mostly related to better strategic decisions. Analytical MDM only should suffice in most cases.

all three business cases described earlier, then obviously, an enterprise MDM implementation is the way to go. For some issues, however, enterprise MDM is simply overkill. This section aims to recommend the right solution to each problem, ensuring that companies do not attempt to kill an ant with a machine gun.

Note that Table 1.1 should be used only as a general guideline. Further, each of the three business cases overlaps the others to an extent; one does not start where the other ends. One could correctly argue that by mitigating certain risks, companies lower the costs of doing business. This, in turn, raises questions of whether this should be in the cost reduction category instead. Also, better data could reduce costs and improve marketing, and consequently grow revenue.

This book will highlight how the MDM practices and techniques described can apply to any of these architectures, or in some cases, may be more relevant to just a certain type of architecture.

DATA MANAGEMENT MATURITY LEVEL

While a business case for MDM will normally address one or more of the three reasons presented earlier, it is important to strengthen it with an accurate

assessment of where the company is regarding the overall MDM spectrum. Two companies with the same business needs will not necessarily follow the same steps to get there because they might be at different maturity levels regarding what is necessary from an MDM practice perspective.

Several data governance maturity models exist and should be used as a frame of reference. Even though data governance in itself is just one of the components in MDM, it is perhaps the most pervasive one since it overlooks all other activities within MDM. That means a data governance maturity model can be used as guidance to understand where a company is and where it should be regarding the management of data as a strategic asset.

Multiple MDM vendors have their own maturity models, but most of the time their models can be used independently from their product lines. Besides, there are models from neutral companies such as Gartner, Data Management Association (DAMA), and the Data Governance Institute (DGI). "Data Governance Part II: Maturity Models—A Path to Progress,"[1] a paper authored by the National Association of State Chief Information Officers (NASCIO), provides a good overview into data governance maturity models, and a closer look into several existing models by multiple vendors.

The end state of a data governance maturity model is essentially when a company achieves a level characterized as being proactively governed, optimized, effective, standardized, and quality controlled at a global level. Everyone in the company will typically agree in principle with the merits associated with achieving this end state or at least with achieving a significant advancement toward this end state. Advancement consists of many individual objectives and achievements, with each one rooted to solving problems that generally everyone should be able to recognize and relate to in the current state. Addressing these problems becomes the heart of the business case. Relating these problems to "here is where we are" in the Maturity Model, and being able to express the achievable objectives that will drive advancement toward the end state, will be extremely important in establishing the business case.

Let's look at some examples of this. The intent here is to provide examples that will express undeniable customer data management problems and their risks along with a realistic course of action needed to mitigate the problem and to create ongoing practices for advancing in the Maturity Model.

These examples obviously need to be tailored to fit the type of issues and objectives that will be relevant to your company. Doing so will greatly help drive a shared perspective about the business problems, strategy, and direction needed to address them. This business case will also serve as a foundation for establishing the MDM model and creating the data governance charter that we

TABLE 1.2 Examples of Customer Data Management Problems

Problem	Risk	Action Needed
There is growing perception in our company that poor data quality is creating customer satisfaction issues, but we don't have any specific data quality measurements to qualify this.	Issues with data integrity, duplication, and fragmentation are going unchecked. This can definitely have an increasing impact on the customer transaction processes.	Data quality metrics and analysis techniques are needed to fully scope this problem in order to determine cause, effect, and a mitigation strategy.
The customer data in our data warehouses is inconsistent and cannot be trusted due to insufficient standards and lack of control at the data-entry level.	Without an accurate source of reference, operational reporting and customer analytics will be subject to error, interpretation, and cross-functional disagreement.	An overarching data governance process is needed to define and drive policies and initiatives that will establish the necessary standards, quality control, and trust of the data.
There is no single view of a customer, and each system has different representations and classifications of the same customer.	A 360° view of the customer cannot be achieved. Marketing, sales, and services cannot effectively synchronize on customer identity and continually are in conflict with customer reporting.	A customer data integration and hierarchy management strategy is needed to build a common customer view to be used as the foundation for consistent business intelligence and customer strategies.
Whenever we discover some major data management or data quality oriented issues, they always seem to require a major escalation followed by an inefficient and time-consuming process to define a focus team just to begin addressing the problem.	Data problems causing major business issues such as related to compliance, privacy, or operational performance, will be continually difficult to resolve and even further exacerbated if there is not a standard process with sufficient resources and roles for data quality management.	As part of a data governance model, an ongoing data quality management forum and process is needed to consistently and quickly address data-oriented issues that can impact business operations.

will cover in more detail throughout this book. The point here is to build a strong central business case that will drive specific, measurable, attainable, realistic, and timely (SMART) goals throughout the MDM initiative.

ADDRESSING THE ROI QUESTION

There is no one recipe for making the case for Customer MDM. Attempts to try to calculate and project ROI will be a swag at best and probably miss the central point that MDM is really an evolving business practice that is necessary to better manage your data, and not a specific project with a specific expectation and time-based outcome that can be calculated up front. Instead, consider all the business dependencies and decisions made that are associated with this data. If anything, the longer-term value of MDM can only be truly measured in real time as cohesive data management and sound governance decisions are made based on ongoing business needs and strategic plans. MDM practices driven by a governance process should certainly consider ROI where possible in making investment and data quality improvement decisions, but MDM, as a developing internal core competency, should be considered as an investment toward improving fundamental data management practices across the company.

Consider what we have identified as the fundamentals of Customer MDM practices. These are process and quality management investment areas that should be justifiable based on any number of existing business problems and data issues. Company A might suffer from such poor data quality that it can barely function. Company B may face severe strict government oversight. In any event, companies first need to recognize the strongest probable benefits of an MDM initiative and build the business case around that. Next, they should estimate how much they are losing by not realizing all the benefits of a having a timely, accurate, and consistent set of data delivered to the company. This is sometimes referred to as activity-based costing (ABC). Often, the best way to measure the potential benefit of MDM involves determining the amount of money that a company spends with reactive activities in place to compensate for a suboptimal set of processes and tools.

SUMMARY

Defining the proper business case is a critical requirement before deciding what Customer MDM approach and architecture to adopt. MDM is not exclusively a

technology issue, but it is also a business capability. As such, companies need to align the proper MDM implementation with their overall strategy to address one or more of the following: risk mitigation, cost reduction, and/or revenue growth.

Not all Customer MDM implementations are the same. Analytical, operational, and enterprise MDM offer different challenges from both IT infrastructure and business model perspectives. Addressing analytical or operational data will lead to distinct levels of complexity, risk, and impact that need to be fully analyzed.

Customer MDM requires bringing many disparate data elements, systems, and processes into a common framework. But this involves a lot more than just data integration. It also requires aligning multiple business units into an integrated set of data, processes, rules, standards, policies, and procedures. It is about fostering collaboration to achieve a high level of success in overall data management, including: data governance and stewardship; data quality; data architecture, analysis, and design; data security; business intelligence; and reference and metadata management.

MDM is *not* a one-time project. It is a very pervasive program that requires executive sponsorship and complete collaboration from the many groups impacted. It also doesn't have a predefined formula. Current data management maturity within the company will dictate the proper steps to be taken, making it very important to position yourself properly from the beginning. Don't be discouraged, though. It is a long and tenuous road but very worthwhile in the long run. If done right, the outputs more than justify the inputs.

NOTE

1. NASCIO, "Data Governance Part II: Maturity Models—A Path to Progress," www.nascio.org/publications/documents/NASCIO-DataGovernancePTII.pdf (NASCIO, 2009).