Managing Master Data with TIBCO Collaborative Information Manager

A Technical Overview

Master data management (MDM) has become a priority for large organizations. Business realities such as global outsourcing, the need for differentiated customer service, better risk management, and more efficient internal operations are forcing organizations to get better insights and value from their information assets. IT strategies to support these initiatives, such as service-oriented architecture (SOA), do not realize their full return on investment when the underlying data they rely on is inaccurate, inconsistent, and often not rich enough to provide business value.

An MDM solution enables businesses to align their enterprise master data assets (product, customer, vendor, etc.) across multiple systems and departments and with trading partners. It also enables organizations to support the necessary processes, policies, and procedures to ensure that accurate and consistent information is disseminated across the organization to the transactional systems and decision-makers that rely on this information for their day-to-day operations.

This paper provides an introduction to the MDM space, the value of an MDM solution to customers, and the pros and cons of common MDM approaches in the marketplace. It discusses TIBCO's MDM solution and unique approach to MDM in the industry. Finally, it provides a technical overview of TIBCO Collaborative Information Manager™ (CIM), TIBCO's MDM application offering, with a focus on the five key capabilities required for a successful MDM deployment.

TIBCO®
The Power of Now®

Introduction

Businesses today must constantly roll-out new value-added services to stay competitive. Retail and manufacturing firms are struggling to meet mandates around collaborative planning, forecasting, and replenishment (CPFR), global data synchronization (GDS), and radio frequency identification (RFID). With mass consolidation in financial services, firms are looking to gain a single customer view for up-sell/cross-sell campaigns. Telecom companies, faced with increasing churn rates and lower switching costs, are focusing on customer service initiatives and better service provisioning. Trading desks across energy, commodities, and securities are trying to better profile and track counterparty risk. Nearly every industry can benefit from better customer service, more efficient procurement, consolidated billing, and more visibility into operations across departments. To support such initiatives IT departments are building strategic composite applications on top of a service-oriented architecture (SOA) backbone. In order for these service-oriented business applications (SOBA), or composite applications, to meet the needs of the business community, IT must achieve some level of information alignment so that different systems being accessed by a business service have consistent definitions and values for critical master data such as product names. Yet most companies do not have accurate and consistent information across people and systems, either internally or across their value chain. This lowers the return on their IT investments.

A master data management (MDM) solution enables businesses to manage and align their enterprise master data assets (product, customer, vendor, etc.) and build and support the necessary processes, policies, and procedures to ensure that clean data stays clean. Master data should not be confused with transactional data. For example, product master data refers to product attributes such as product name, description, dimensions, and other such attributes. Individual order information such as PO numbers, order quantity, etc. are examples of transactional data. A custom attribute such as credit score, which might be computed using transactional data and updated more frequently than a static attribute such as customer name, is still considered and managed as master data, as it is a descriptor of who that customer is, similar to 'customer address' or 'income bracket'. Managing such a rich and complex array of ever-changing and interdependent attributes across multiple data domains on an ongoing basis and then ensuring that this data filters down across the multitude of transactional systems and decision-makers that depend on it for their day-to-day operations is the power of an MDM solution.

Aligning master data internally across multiple instances of enterprise resource planning (ERP), customer relationship management (CRM), and supplier relationship management

(SRM) systems, to name a few, is only a first step, albeit an important step. Just as important is making sure that information – such as product names and attributes or customer data – is consistent across the value chain so that companies can reap the benefits of stronger value chain collaboration such as smaller inventory buffers in a supply chain, better customer service for large customers, and reduced order-to-cash cycles through more efficient purchasing and invoicing.

Why Do Customers Need MDM?

Master data management refers to the generic problem of managing enterprise master data assets and empowering downstream people, systems, and partners with the information they need. Different organizations may choose to focus on a different area or domain of master data. Nearly every company can benefit from more accurate information on their products, customers, and vendors across industries. Common challenges that MDM addresses include:

- A large number of products, parts, or assets, each with multiple attributes that need to be managed and enriched by different departments
- Need to gain more insights into customers for both new revenue opportunities as well as better customer service across multiple customer interaction points and channels
- Need to harmonize inconsistencies across systems, business units, and geographies
- Need to execute business processes such as new client introduction that cut across information silos and organizational boundaries – geographic and functional
- Need to receive client and security information from external feeds such as Reuters and Bloomberg for trade settlement
- Need to collaborate with other businesses in the value chain to achieve efficiency
- Partner and regulatory compliance mandates requiring the management of new types of information and the maintenance of data lineage trails
- Need to measure metrics that require aggregation across multiple systems such as counterparty risk across all trading assets
- Need to gain efficiencies in procurement across multiple products and vendors

Most companies will choose to tackle one data domain first and then use that experience to expand into other domains. Many industries will have data domains that are uniquely valuable to them. For example retail chains may want to first focus on store information.

An auto manufacturer may want to manage dealer information. Healthcare companies will have unique challenges around patient information. While these different data domains may appear very different to a business analyst, from a technology standpoint they pose similar challenges.

- **Comprehensive Information Management:** Managing the data model and attribute information, validation and transformation rules, versioning and diffanalysis, roles-based access and ownership control, complex data relationships including management across data domains, classifications, contextual validation rules, etc.
- **Process Management:** Managing the processes and procedures around introducing new data or editing existing data such as launching a new product or updating a customer address.
- **Integration:** Synchronizing in real-time or batch the relevant subset of master information with transactional systems such as ERP and trading partners either directly or through exchanges and data pools such as 1Synch.

MDM DRIVES ROI FOR SOA INVESTMENTS

An MDM solution provides the necessary alignment of master data across multiple back-end systems so that business services and composite applications within an SOA have accurate, consistent, and timely information. All the hype and attention in SOA deployments has gone into web service creation, deployment, and management standards and technologies. However, if data is inconsistent across applications it will be increasingly difficult if not prohibitive to build composite applications that cut across multiple systems and departments. For example a composite application in a large multi-channel financial service institution that calculates a customer's global credit risk will only work if that customer is described in a consistent manner across retail banking, brokerage, mortgage, and credit card systems. In a retail environment, a composite application that gets a customer's order history from a data warehouse and recommends a related product will require consistent product and customer information across all the relevant systems. On a smaller scale, even business services to update an address or provision a service require semantic consistency of master data across CRM, billing, and product systems. Creating a semantic integration layer to harmonize and reconcile master data across the enterprise yields accurate and consistent information that helps SOA investments realize their full ROI.

Common Approaches to Managing Master Data

Today there are three competing approaches to managing master data in the marketplace:

- An independent harmonized master data solution. This approach builds and maintains a central information repository and then synchronizes that repository with back-end transactional systems through a publish-subscribe API.
- Storing master data in a dominant application such as the ERP or CRM system. All other systems that need access to that information treat that system as the system of record.
- A federated approach where data is not maintained in a central repository. In this approach a cross-reference registry of global-IDs is stored and used to access and aggregate information across systems on-the-fly in a virtual manner. This approach is also referred to as enterprise information integration (EII).

DOMINANT APPLICATIONS

The second approach, treating the ERP or CRM system as the system of record, is perhaps the most common approach in the marketplace today. The driver for this has largely been inertia rather than a sound technical reason. Firms have made huge multimillion dollar investments in these large applications and are using what they have bought regardless of the applicability to the problem – trying to extract the most value for their buck. The problem with this approach is:

- It is very rare for large companies to have just one ERP system. Most large organizations, either through organic growth or acquisitions, have multiple ERP systems and multiple instances of each ERP across business units and geographies. Designating any one of these systems as the system of record becomes both politically and technically impractical.
- The inherently siloed nature of these applications makes it very difficult to develop cross-domain relationships such as between customers and products or between products, vendors, and store locations.
- Dominant applications are inherently transactional in nature and not designed to manage the large superset of data attributes that make up master data. They are focused on the attributes relevant to them, which are not necessarily all the attributes about that product or customer that need to be managed to satisfy enterprise needs.

- ERP and CRM applications typically provide minimal business process configurability outside of transaction processing. MDM itself requires a collaborative approach to introduce and update data in a controlled manner that conforms to corporate data governance standards.
- Finally, these systems are fairly rigid by design in terms of their data structure and business processes to protect the integrity of their transactions. Custom development to infuse process and data model extensibility is time consuming, expensive, risky, and often does not deliver the intended business benefits.

FEDERATED

The federated or EII approach has gotten some traction as a new way to tackle the problem. Its upside is that it is not plagued by the political battles around data ownership as it leaves data where it lies and accesses it dynamically. While this approach may seem advantageous at first glance, it has significant drawbacks.

- Access only: It is far easier to access data from many distributed sources than make an update. If many systems have to be updated to complete a transaction, the federated solution will have to build complex handshaking and roll-back logic to be viable.
- Only good for small amounts of data: The federated model is only suitable for small sized information movements, as it has to access many systems in real time and perform in-memory aggregations, filtering, and transformation on data that may not be optimized to be viewed in the manner being requested.
- Most importantly, the federated model does not fix the data problem; it works around it. If the organization is plagued by inaccurate, redundant, and inconsistent master data, the federated approach to MDM does not fix the underlying data; it merely puts a favorable view on top of the data. Data quality is only as good as the federated data sources, resulting in the same political issues as it was designed to avoid.
- Information management functionality such as versioning, rollback, audit control, or validation is handled by the individual systems and can be inconsistent or inadequate. Furthermore, validating data relationships for data stored across multiple systems is not possible.

Figure 1. Common approches to MDM

HARMONIZED	FEDERATED		WITHIN DOMINANT APPLICATION
Data is Persisted	Yes	No	Yes
Multi -Domain	Yes	Yes	No
Information Management Prcesses	Supported, Customizable	Not directly supported	Supported but fixed
Data Usage	Read/Write. Serves asa central reference point but not used intransactions.	Generally read only	Read/Write, Transactional
Information Sources	Heterogeneous	Heterogeneous	
Synchronization Frequency	Real-time andbatch	N/A	Real-time and batch
Data Model	Flexible	Flexible	Generally fixed

HARMONIZED – TIBCO'S APPROACH

Given the pitfalls in storing master data in a dominant application or having a completely federated approach, TIBCO believes that the harmonized model for MDM yields the most value for customers. TIBCO Collaborative Information Manager software has the following attributes that make it unique in the marketplace.

Harmonized: TIBCO builds and maintains a central master data repository that serves as either the system of record or reference and can be a source of authoring of information or designate authoring to transactional systems. As the system of reference it keeps track of ownership at the attribute level and applies precedence rules to ensure survivorship is accurately handled when importing new or changed data. The repository then synchronizes with multiple downstream transactional systems and trading partners that depend on that data. The master data repository can also serve master data as services to requesting composite processes over a Web services API. The repository is populated by aggregating, cleaning, transforming, and enriching data from internal and external sources and maintained by the process-centric capabilities of the application.

Process-Centric Approach: TIBCO takes a process-centric approach to information management by providing customizable processes to introduce and manage data, such as a "new product introduction" process in a manufacturing company or a "new vendor

introduction" process in telecom. If data is cleansed and aggregated, such as in data warehousing projects but the necessary business process automation and data governance policies and controls to introduce and update new data are not put in place, the value of the solution will be fleeting and over time the original problem of data inaccuracies and inconsistencies will return.

Collaborative: The process-centric nature of the application allows it to be extremely collaborative. Different attributes of data will always be owned and managed by different business users. For example, product brand data will be managed by marketing, packaging data by logistics, and product specs by R&D. TIBCO CIM facilitates cross-department and cross-geography collaboration to manage, update, and enrich master data.

Multi-Domain: TIBCO provides a single application to manage data across multiple strategic data domains, be it customer, product, vendor, etc. It does so through an extensible configurable data model and customizable information management processes to meet the unique needs of different types of information. Managing multiple domains on a single platform leads to a low TCO and allows customers to scale their solution as their needs expand. In addition, customers can also manage relationships between data domains such as between products and vendors, or securities and clients.

SOA: TIBCO has over 20 years of experience in distributed computing and has pioneered publish-subscribe technologies. TIBCO leverages these strengths to provide a highly distributed solution that can scale to meet the demands of very large deployments.

Figure 2. MDM in a manufacturing environment

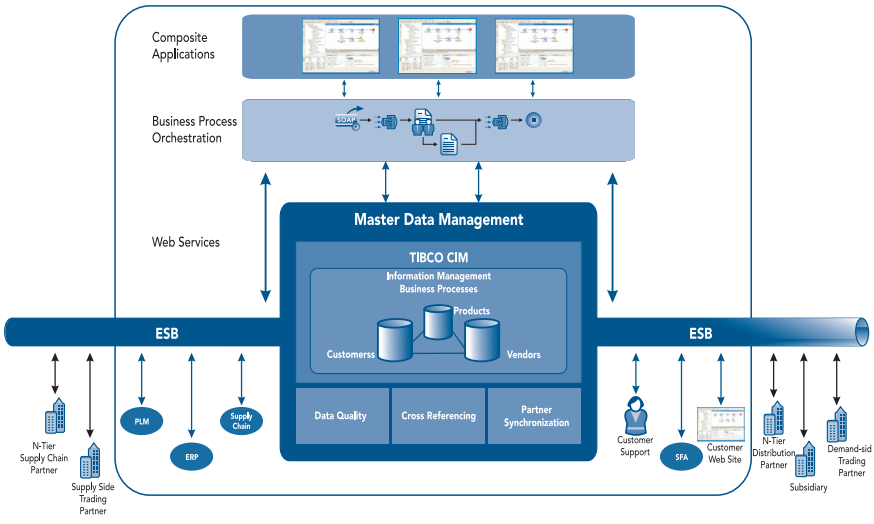

The TIBCO CIM repository is synchronized via both batch and real-time technologies with internal and partner systems. Through features such as the distributed cache, Web services API, support for multi-threaded asynchronous processing, event-based architecture, and embedded JMS client, TIBCO CIM is the leading solution for customers with SOA initiatives or looking to support SOA principals for the long term.

Peer-to-Peer: TIBCO CIM supports peer-to-peer synchronization with other TIBCO CIM instances. Such an architecture might be necessary in a global organization where regional or business unit differences are so great that it makes sense to keep a local master data repository in each region or departnment and synchronize it with corporate headquarters and other regions. This is common in multi-banner retailers, or large conglomerates looking to provide regional autonomy with centralized visibility and control.

A Complete MDM Solution

For an MDM solution to be successful it has to be tightly integrated in a bi-directional manner with the underlying IT landscape as well as trading partner systems, which are out of your control. TIBCO CIM supports multi-channel bi-directional state-full synchronization with internal systems, trading partners, and industry exchanges such as 1Synch.

The solution must be flexible to account for the fact that the link with every system will be governed by its own set of data transformation and validation rules, synchronization frequency (batch, near real-time, real-time), and integration transport (FTP, point-to-point JMS, web services, message bus, AS/2).

During the on-ramping of the solution, data needs to be profiled, cleansed, aggregated, and loaded in bulk with validation rules into the repository. After deployment, the entire solution, including the core TIBCO CIM application and all data feeds going in and out of it, needs to be monitored in a mission-critical manner.

Finally, ongoing data management processes and key metrics have to be tracked for effectiveness for continuous process improvement.

TIBCO provides a complete MDM solution with a high-powered business integration platform encompassing enterprise application integration (EAI), business-to-business (B2B), enterprise service bus (ESB), and extraction, transformation and loading (ETL) capabilities as well as necessary monitoring and management. All TIBCO products, while designed to integrate well with each other, are open and standards-based. Hence, the application can work with whatever integration platforms customers already have in-house. While TIBCO

provides all the supporting components mentioned above, no other TIBCO products are required other than TIBCO CIM for MDM. Customer who wish to leverage their existing middleware infrastructure can do so. TIBCO CIM can be deployed as a stand-alone MDM application.

Figure 3. TIBCO Collaborative Information Manager architecture

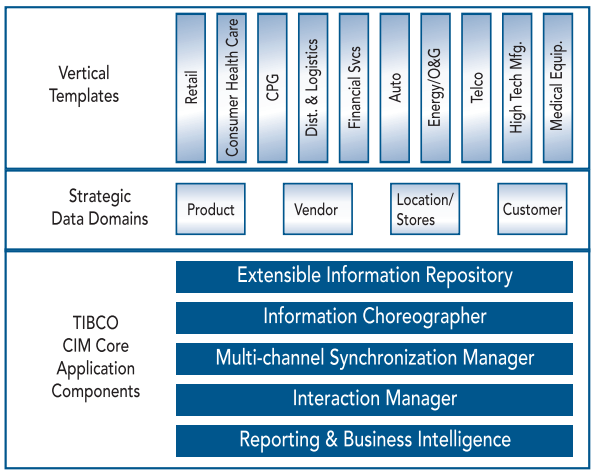

TIBCO Collaborative Information Manager Architecture

The TIBCO Collaborative Information Manager application is the foundation of TIBCO's MDM solution. TIBCO CIM delivers the key components necessary to achieve comprehensive MDM.

- Core application components
  - Information repository
  - Data governance
  - Process Automation
  - Internal and external synchronization
  - Reporting and business intelligence
- Support for multiple strategic data domains
  (product, customer, vendor, location, financial data, reference data, code tables, etc.)
- Vertical templates

LOOSELY COUPLED PRODUCT ARCHITECTURE

The TIBCO CIM application is designed in a loosely coupled manner with well defined standards-based interfaces between its internal components. Process templates are built on top of a flexible data model that relies on the five core application components to power the solution. Abstracting these layers allows for a highly flexible and scalable offering that can be quickly configured to meet the customer's requirements. Furthermore it provides the ability to support multiple data domains (product, customer, vendor, etc.) from one platform.

GUI

CIM is an entirely thin client application except for a thick-client optional graphical process modeler. The web user interface is entirely metadata driven so screens are dynamic and contextual. Information is passed back and forth between the engine and UI components as XML; hence it is easy to swap out TIBCO CIM screens and embed custom rich-client interfaces right into the application using tools such as TIBCO's Ajax-based product TIBCO General Interface™.

The CIM GUI is based on the user interface design practices as laid out by the TIBCO One initiative, a corporate wide initiative to provide customers with a best-in-class unified UI experience across products.

The CIM GUI uses web 2.0 technologies such as AJAX throughout the application for graphically creating input maps, installation and configuration, building relationships, etc.

High Performance

TIBCO CIM makes intelligent use of distributed caching and parallelization of work across multiple threads to load, process, and publish large amounts of information very quickly. This is essential in data-intensive industries such as financial services, telecom, and energy. TIBCO CIM can automatically split large record sets into smaller batches for asynchronous processing with parallel threads working on individual batches. This greatly enhances performance and allows for the use of all available hardware by the system. Batch sizes can be configured via a property file. Inserts into the repository also occur in larger batch sizes rather than as individual transactions. Additionally, TIBCO CIM can be configured to load record data into the cache at startup. The validation and process engines would then be retrieving information from the cache and could evaluate the rule base and load the data into the repository considerably faster than before.

TIBCO CIM aggressively uses the cache to speed performance by caching workflow states, Web service responses, user login information, temporary objects, etc. By gradually moving to a cache-first architecture with intelligent refresh and write-behind CIM leads the market

with respect to performance and high volume deployments.

The TIBCO professional service group provides best practice guidance on making optimal performance decisions such as structuring the data model and relationships, instrumenting activities in the workflow, and executing activities in a synchronous or asynchronous manner

WEB SERVICES

Application functionality – down to a fairly granular level such as querying and updating the structure of the master catalog, adding or modifying records, firing a business process, or querying and submitting work items within a business process, to name a few – are all exposed via pre-packaged web services. This makes it easy for customers who wish to use their own custom UI outside the application to call TIBCO CIM services or embed TIBCO CIM services as part of a larger composite application. Both functional and highly granular web services are provided out of the box, so customers have a great deal of flexibility in accessing and controlling TIBCO CIM functionality from outside the application. This yields a master data services platform to support SOA and BPM initiatives that need reliable access to accurate and consistent data.

Application-Based: TIBCO CIM is a self-contained MDM application that provides a tremendous amount of out-of-the-box functionality, including data models, rules, and business processes. TIBCO CIM has a wizard-driven installation and configuration process and does not require any other TIBCO software and can work with other vendor's integration and business-to-business platforms as necessary.

Flexibility: By embedding the rules engine, process engine, and synchronization engine as part of the application, these engines are "aware" of the data model. This means that any changes to the data model such as adding or modifying an attribute are instantly available to the rules and process engine and instantly available for downstream consumption without the need to re-code any interfaces. Further, changes to the data model do not disrupt the system and can be hot-deployed. This versatility differentiates the application from the competition and lowers the cost of ownership.

Configuration vs. Customization: A solution that can be configured rather than built at the client site leads to faster deployments and a lower TCO. This includes configuring the data model via the GUI, importing existing metadata via the GUI, configuring rules either via the GUI or through the XML source files, and configuring processes via a graphical process modeler without the need to write code. The configuration vs. customization approach leads to seamless upgrades, lower services costs, and higher project success rates.

STANDARDS

TIBCO CIM is a completely standards-based Java EE web application with a zero-client footprint. It is event-driven, providing a great deal of flexibility and control over information management processes. Access rights and ownership of data is role-based and can leverage existing LDAP systems. Business processes and associated rules are stored as XML and can be modified via the application GUI or directly by a developer. Communication with the outside world is done over standard protocols such as FTP, JDBC, web services, XML over JMS, or AS/2. Monitoring can be done by a standard JMX console. No specialized or proprietary skill sets are required for the implementation or management of TIBCO CIM.

Five Critical Components of Collaborative Information Manager

TIBCO Collaborative Information Manager is comprised of five key capabilities that power the application:

- Information repository
- Data governance
- Process Automation
- Internal and external synchronization
- Reporting and business intelligence

INFORMATION REPOSITORY

TIBCO CIM offers comprehensive information management through an extensible and flexible data model with support for multiple classification schemas, product and attribute relationships, and automated data validation.

Data Model: The TIBCO CIM data model is designed to be extensible to meet the customer's needs. Metadata can be imported and exported to leverage existing modeling work the customer may already have done. With an extensible data model, a business analyst or power user can expand and customize the data attributes from the GUI without the need for a DBA or data architect. For example, a securities master may have a rich complex series of attributes depending on whether the instrument is a stock, corporate bond, or treasury security and these attributes could further vary across regions. A supermarket chain may want a Boolean attribute to be called "sugar-free" for a candy bar, but the same attribute to be called "non-fat" for a different product (or the same product

in a different region). Unstructured data such as images, R&D specs, bar codes, and presentations can also be stored, managed, and associated with records. To truly serve the needs of the organization the data model must be easily customizable and configurable to accommodate a wide variety of data types, attributes, and relationships. For an organization that depends on large sets of master data, accurately managing this rich superset of attributes and relationships in a consistent manner across the enterprise can be next to impossible without an MDM solution.

Relationships: TIBCO CIM enables the definition of relationships between records such as peer-to-peer to parent-child relationships. For example, a telecom vendor can define relationships across products for local service, long distance, and DSL, thereby taking advantage of up-sell and cross-sell opportunities. In retail these relationships can be dynamic and complex and may change seasonally or regionally, accounting not only for complementary items but substitutions as well. These relationships can also be hierarchical in nature, as between a client and subsidiary. Relationships can have their own rules enforced by configurable validations and dependencies, such as every "case" item must have 12 "can" items.

Cross-Domain Relationships: The power of being able to manage multiple domains of information from a single platform is truly realized by the ability to create cross-domain relationships based upon rules. For example, a vendor record in a vendor catalog can have a relationship to the products that the vendor sells to the company. Or a product could be tied to a customer, region, or promotion. For many organizations the relationships between their data domains is almost as valuable as the data within the domain itself. For example for an organization to optimize their supply chain they need to know where products are sourced from and which customers they typically get shipped to. TIBCO CIM provides powerful cross-domain features. Users can easily model what they want these relationships to look like from the GUI. Validations can be performed across these relationships such as ensuring that only approved vendors are being selected while setting up a new product. Importing data can be done from either a single feed or multiple sources and then inserted into multiple catalogs with the right relationships in place. Similarly out-of-the-box web services allow for information across repositories to be queried or inserted as a single bundle from an external application.

Classifications: TTIBCO CIM supports the creation and maintenance of unlimited classifications schemes and provides pre-built classification schemes where industry standards exist, such as UDEX in retail. Business users can also define their own classification schemes based on what makes the most sense for them, such as 'income bracket' for customers.

Search: As these repositories grow to the millions and tens of millions of records a powerful and flexible search capability is essential to help business users find what they need in an efficient manner. TIBCO provides both structured and context-free Google-like search capabilities that can span multiple catalogs at once either from the GUI of a Web service. Fuzzy matching service is also provided via a Web service that can accept attributes weights as parameters and then returns a match score.

Versioning: Both data and metadata are versioned and a detailed audit trail is maintained. If the attributes associated with an record change over time, these changes can be tracked, compared, and rolled-back. Different versions of data and metadata can be compared to allow for questions such as, "What did this vendor address look like 6 months ago?" This is also essential to support long-lived processes, such as a new account introduction that might take days or even weeks to complete. Version comparisons can also be process-based. For example, when a data attribute changes, the application can automatically compare it to the previous version and send an alert if this change is greater than a specificed threshold. The last confirmed version is typically referred to as the "golden copy." TIBCO, while storing rich version history, also provides a separate table and easy mechanism for external processes and applications to query just the golden copy and not have to sift through unconfirmed or previous versions. Customers can purge and archive the data base to keep the repository size manageable.

DATA GOVERNANCE

Validation: TIBCO CIM provides the ability to automate all validation rules for data both entering and leaving the master data repository. For example, a user could create a rule to ensure that currency entered is always in Dollars or Euros. Validation can also be contextual or rules-based, for example, ensuring that the system knows that net weight of a product must always be lower than gross weight or that certain services or products, such as DSL service, are only supported in certain geographies. Validation rules are abstracted from the data model for flexibility and can be added at any time. Valid value lists for attributes can be maintained inside the application or linked to dynamically. Attributes can be declared as optional or mandatory. By providing out-of-the-box and configurable validation rules for all new data, the application ensures the integrity of all information within the repository.

Security and Access Control: TIBCO CIM allows for securing any resource within the application and assigning a different level of access to any role or user. A resource can be any object managed within the system, be it a record, an attribute, a catalog, catalog subset, or process. In certain cases, records may need to be secured based on values contained

within the attributes such as restrict access to all customer records where net worth is greater than $1M.. Single-sign on is supported as well as synchronization wit LDAP.

ADVANCED BUSINESS PROCESS AUTOMATION

TIBCO CIM provides the ability to automate master data creation and maintenance processes such as introducing a new product or product bundle, launching a new distribution channel, adding or removing a vendor or customer, etc. This puts procedural rigor around data management so that the act of aggregating, cleansing, and enriching data is not a one-time exercise but an ongoing collaborative mindset that filters across the organization.

Rich Out-of-the-Box Functionality: The workflow engine within TIBCO CIM is extremely rich and gives users a great deal of flexibility and ease of use to create and modify processes.Support for long-running processes, ability to execute activities serially or in parallel, deadline-driven reminders, versioning and rolling back of processes, escalation and reassignment of work due to delays or vacation time, multiple users to a role with work item locking, iterative rejection approval loops, and detailed audit trails and process monitoring are just some of the features built into the product to help make organizations more productive.

Graphical Process Modeler: TIBCO provides a complimentary product called TIBCO CIM Process Designer (CPD) which is a graphical process modeler for CIM workflows. TIBCO CPD allows technical users to drag and drop activities and transitions and configure tasks to create rich workflows. TIBCO CPD is based on TIBCO Business Studio, TIBCO's eclipse-based open-source common workflow modeling environment.

Collaborative: Even though data is replicated centrally within the CIM repository, its ownership can be distributed. Different business users across functional roles are responsible for creating and maintaining different attributes. TIBCO CIM enables this through role-based access rights. For example, marketing may own the product description but finance may own sales tax information

Business-User Friendly: Empowering business users with accountability and ownership of relevant data is further enhanced by the fact that the UI screens used to add, edit, and delete data elements are very business-user friendly and do not require an administrator. To some extent, even the creation and modification of business rules and processes can be done by business users, without assistance from an administrator. By putting ownership in the hands of the natural owners of the data, the business users themselves, the solution takes much of the burden of data management off of IT. In fact it is TIBCO's experience that

until the business users get involved and commit to doing this, it is very hard to make an MDM initiative successful. MDM must be a collaborative endeavor across business users, IT, and senior management.

INTERNAL AND EXTERNAL SYNCHRONIZATION

TIBCO CIM leverages TIBCO's strengths in integration to manage information synchronization with back-end systems and partners using standard technologies and protocols. TIBCO CIM supports multiple integration transports, from simple file-based FTP to sophisticated enterprise service bus (ESB) technology, depending on the complexity of the customer's environment. TIBCO CIM also provides pre-built transformation maps that support industry message standards such as EAN.UCC in retail. The open, standards-based distributed nature of the application facilitates robust integration with TIBCO's industry-leading business integration platform or any of the major integration platforms in the marketplace.

Process-based: Synchronization is not fire and forget or even pub-sub. Instead an entire synchronization can be orchestrated per target to ensure the necessary level of handshaking is executed to ensure reliable delivery and exception management.

Event-driven: TIBCO CIM's event-driven architecture allows the application to call out to external applications during workflow processes and take action or spawn new processes based on the results, such as fetching vendor information from a vendor management system during a new product setup process.

Target-driven: TIBCO CIM provides the capabilities to create publications. A publication is a configuration that specifies what subset of data is sent to which target over what transport and on what schedule. TIBCO CIM does not publish all master data at once and burden the network to parcel subsets of information to different targets. It also supports the ability to only publish information that has changed.

Standards-based: TIBCO CIM is platform neutral for integration and follows standards such as XML, JMS, Web services and AS/2. This allows CIM to leverage TIBCO's middleware products or whichever standards-based middleware products a customer already has in-house.

Cache: TIBCO CIM embeds a distributed cache to serve up data and metadata to requesting applications in a service oriented manner with high throughput. This allows CIM to serve as a very scalable master data services platform.

Partner Synchronization: In today's world of global outsourcing and complex value chains, nearly every organization must obtain and synchronize information with trading partners. TIBCO CIM provides automatic synchronization with trading partners and some industry exchanges such as 1Sync (UCCnet, Transora) and Agentrics (WWRE) in retail. TIBCO CIM handles the message management and choreography while keeping track of state. Partners only receive the information they care about in the format of their choosing. Integration outside of the firewall is done via AS/2.

REPORTING AND BUSINESS INTELLIGENCE

Organizations need insight into their information assets and the processes that manage them to make informed and timely business decisions. Often by the time data makes it into a data warehouse for analysis it is already too late. Furthermore, this data is static and may not have business process context. TIBCO CIM enables real-time measurement of key performance indicators (KPIs) so that managers can make opportune choices to help ensure corporate health. For example, managers may wish to measure sales volumes or profitability by vendor in real time, but operational systems do not have the cross-product, cross-customer, cross-region relationship links and classification schemes to effectively rollup this information into aggregates. By storing and managing these relationships, TIBCO CIM helps make existing information more useful.

TIBCO CIM also supports the measurement and effectiveness of information management processes over time such as new customer introduction process so that bottlenecks can be identified and these processes and be continuously tuned. TIBCO CIM stores detailed event, data and metadata information that can feed enterprise reporting toolsas well as captures a large array of statistics about data usage.. The richness of the data lineage maintained by TIBCO CIM as well as features such as compare and roll-back can give valuable insights into the health of master data and support compliance initiatives.

For more information, visit http://tibco.com.

Conclusion

Organizations today continue to innovate at a rapid pace to drive competitive advantage. IT has made strides in areas such as service-oriented architectures and composite applications to quickly roll out new services to support business needs. However, for these investments to be productive it is imperative for organizations to align their master data assets (products, customers, vendors, etc.) across their internal transactional systems and with trading partners.

TIBCO provides a comprehensive process-centric master data management solution. TIBCO's MDM solution provides a central repository and point of reference for master data outside of dominant applications like the ERP system and then synchronizes relevant subsets of master data with all downstream transactional applications, trading partners, and industry exchanges in the right format. Furthermore, TIBCO's solution provides the necessary processes and rules to manage the introduction of new data so that the clean data stays clean. TIBCO's solution is multi-domain, providing a single platform for all types of master data, be it customer, product, vendor, location or financial.

TIBCO CIM, the cornerstone of TIBCO's MDM solution, is a self-contained application that provides rich out-of-the-box functionality to deliver the key capabilities for a master data management solution, including an extensible repository, data governance enforcement, advanced business process automation, internal and external synchronization, and business intelligence and reporting. Industry templates and horizontal best practices for the data model, validation rules, transformation maps, and information management processes are embedded directly into the application for rapid deployment and a low TCO.

TIBCO augments the Collaborative Information Manager application with all the necessary components for a complete MDM solution, including application and data integration, a B2B gateway, data quality and cleansing, Web service support, message-oriented middleware, and robust monitoring and management. With over 20 years of experience in distributed computing for global organizations across industries, TIBCO delivers a best-in-class MDM solution.

TIBCO®
The Power of Now®

Global Headquarters
3303 Hillview Avenue
Palo Alto, CA 94304

Tel: +1 650-846-1000
Toll Free: 1 800-420-8450
Fax: +1 650-846-1005

www.tibco.com