7

Operational Intelligence: Technological Components

Operational intelligence has a number of aspects that have been elucidated in this chapter. Some of these features are complex event processing, business process management, metadata and root cause analysis. The components discussed in this text are of great importance to broaden the existing knowledge on operational intelligence.

Operational Intelligence

Operational intelligence (OI) is a category of real-time dynamic, business analytics that delivers visibility and insight into data, streaming events and business operations. Operational Intelligence solutions run queries against streaming data feeds and event data to deliver real-time analytic results as operational instructions. Operational Intelligence provides organizations the ability to make decisions and immediately act on these analytic insights, through manual or automated actions.

Purpose

The purpose of OI is to monitor business activities and identify and detect situations relating to inefficiencies, opportunities, and threats and provide operational solutions. Some definitions define operational intelligence an event-centric approach to delivering information that empowers people to make better decisions.

In addition, these metrics act as the starting point for further analysis (drilling down into details, performing root cause analysis — tying anomalies to specific transactions and of the business activity).

Sophisticated OI systems also provide the ability to associate metadata with metrics, process steps, channels, etc. With this, it becomes easy to get related information, e.g., “retrieve the contact information of the person that manages the application that executed the step in the business transaction that took 60% more time than the norm,” or “view the acceptance/rejection trend for the customer who was denied approval in this transaction,” or “Launch the application that this process step interacted with.”

Features

Different operational intelligence solutions may use many different technologies and be implemented in different ways. This section lists the common features of an operational intelligence solution:

- Real-time monitoring
- Real-time situation detection
- Real-time dashboards for different user roles
- Correlation of events
- Industry-specific dashboards
- Multidimensional analysis
    - Root cause analysis
    - Time Series and trend analysis
- Big Data Analytics: Operational Intelligence is well suited to address the inherent challenges of Big Data. Operational Intelligence continuously monitors and analyzes the variety of high velocity, high volume Big Data sources. Often performed in memory, OI platforms and solutions then present the incremental calculations and changes, in real-time, to the end-user.

Technology Components

Operational intelligence solutions share many features, and therefore many also share technology components. This is a list of some of the commonly found technology components, and the features they enable:

- Business activity monitoring (BAM) - Dashboard customization and personalization
- Complex event processing (CEP) - Advanced, continuous analysis of real-time information and historical data
- Business process management (BPM) - To perform model-driven execution of policies and processes defined as Business Process Model and Notation (BPMN) models
- Metadata framework to model and link events to resources
- Multi-channel publishing and notification
- Dimensional database
- Root cause analysis
- Multi-protocol event collection

Operational intelligence is a relatively new market segment (compared to the more mature business intelligence and business process management segments). In addition to companies that produce dedicated and focussed products in this area, there are numerous companies in adjacent areas that provide solutions with some OI components.

Operational intelligence places complete information at one's fingertips, enabling one to make smarter decisions in time to maximize impact. By correlating a wide variety of events and data from both streaming feeds and historical data silos, operational intelligence helps organizations gain real-time visibility of information, in context, through advanced dashboards, real-time insight into business performance, health and status so that immediate action based on business policies and processes can be taken. Operational intelligence applies the benefits of real-time analytics, alerts, and actions to a broad spectrum of use cases across and beyond the enterprise.

One specific technology segment is AIDC (Automatic Identification and Data Capture) represented by barcodes, RFID and voice recognition.

Comparison with Other Technologies or Solutions

Business Intelligence

OI is often linked to or compared with business intelligence (BI) or real time business intelligence, in the sense that both help make sense out of large amounts of information. But there are some basic differences: OI is primarily activity-centric, whereas BI is primarily data-centric. As with most technologies, each of these could be sub-optimally coerced to perform the other's task. OI is, by definition, real-time, unlike BI or "On-Demand" BI, which are traditionally after-the-fact and report-based approaches to identifying patterns. Real-time BI (i.e., On-Demand BI) relies on the database as the sole source of events.

OI provides continuous, real-time analytics on data at rest and data in-flight, whereas BI typically looks only at historical data at rest. OI and BI can be complementary. OI is best used for short-term planning, such as deciding on the "next best action," while BI is best used for longer-term planning (over the next days to weeks). BI requires a more reactive approach, often reacting to events that have already taken place.

If all that is needed is a glimpse at historical performance over a very specific period of time, existing BI solutions should meet the requirement. However, historical data needs to be analyzed with events that are happening now, or to reduce the time between when intelligence is received and when action is taken, then Operational Intelligence is the more appropriate approach.

Systems Management

System Management mainly refers to the availability and capability monitoring of IT infrastructure. Availability monitoring refers to monitoring the status of IT infrastructure components such as servers, routers, networks, etc. This usually entails pinging or polling the component and waiting to receive a response. Capability monitoring usually refers to synthetic transactions where user activity is mimicked by a special software program, and the responses received are checked for correctness.

Complex Event Processing

There is a strong relationship between complex event processing companies and operational intelligence, especially since CEP is regarded by many OI companies as a core component of their OI solutions. CEP companies tend to focus solely on development of a CEP framework for other companies to use within their organisations as a pure CEP engine.

Business Activity Monitoring

Business activity monitoring (BAM) is software that aids in monitoring of business processes, as those processes are implemented in computer systems. BAM is an enterprise solution primarily intended to provide a real-time summary of business processes to operations managers and upper management. The main difference between BAM and OI appears to be in the implementation details — real-time situation detection appears in BAM and OI and is often implemented using CEP. Furthermore, BAM focuses on high-level process models whereas OI instead relies on correlation to infer a relationship between different events.

Business Process Management

A business process management suite is the runtime environment where one can perform model-driven execution of policies and processes defined as BPMN models. As part of an operational intelligence suite, a BPM suite can provide the capability to define and manage policies across the enterprise, apply the policies to events, and then take action according to the predefined policies. A BPM suite also provides the capability to define policies as if/then statements and apply them to events.

Business Activity Monitoring

Business activity monitoring (BAM) is software that aids in monitoring of business activities, as those activities are implemented in computer systems.

The term was originally coined by analysts at Gartner, Inc. and refers to the aggregation, analysis, and presentation of real-time information about activities inside organizations and involving customers and partners. A business activity can either be a business process that is orchestrated by business process management (BPM) software, or a business process that is a series of activities spanning multiple systems and applications. BAM is an enterprise solution primarily intended to provide a real-time summary of business activities to operations managers and upper management.

Goals and Benefits

The goals of business activity monitoring is to provide real time information about the status and results of various operations, processes, and transactions. The main benefits of BAM are to enable an enterprise to make better informed business decisions, quickly address problem areas, and re-position organizations to take full advantage of emerging opportunities.

Key Features

One of the most visible features of BAM solutions is the presentation of information on dashboards containing the key performance indicators (KPIs) used to provide assurance and visibility of activity and performance. This information is used by technical and business operations to provide visibility, measurement, and assurance of key business activities. It is also exploited by event correlation to detect and warn of impending problems.

Although BAM systems usually use a computer dashboard display to present data, BAM is distinct from the dashboards used by business intelligence (BI) insofar as events are processed in real-time or near real-time and pushed to the dashboard in BAM systems, whereas BI dashboards refresh at predetermined intervals by polling or querying databases. Depending on the refresh interval selected, BAM and BI dashboards can be similar or vary considerably.

Some BAM solutions additionally provide trouble notification functions, which allows them to interact automatically with the issue tracking system. For example, whole groups of people can be sent e-mails, voice or text messages, according to the nature of the problem. Automated problem solving, where feasible, can correct and restart failed processes.

Deployment Effort

In nearly all BAM deployments extensive tailoring to specific enterprises is required. Many BAM solutions seek to reduce extensive customization and may offer templates that are written to solve common problems in specific sectors, for example banking, manufacturing, and stockbroking. Due to the high degree of system integration required for initial deployment, many enterprises use experts that specialize in BAM to implement solutions.

BAM is now considered a critical component of Operational Intelligence (OI) solutions to deliver visibility into business operations. Multiple sources of data can be combined from different organizational silos to provide a common operating picture that uses current information. Wherever real-time insight has the greatest value, OI solutions can be applied to deliver the needed information.

Processing Events

All BAM solutions process events. While most of the first BAM solutions were closely linked to BPM solutions and therefore processed events emitted as the process was being orchestrated, this had the disadvantage of requiring enterprises to invest in BPM before being able to acquire and use BAM. The newer generation of BAM solutions are based on complex event processing (CEP) technology, and can process high volumes of underlying technical events to derive higher level business events, therefore reducing the dependency on BPM, and providing BAM to a wider audience of customers.

Examples

A bank might be interested in minimizing the amount of money it borrows overnight from a central bank. Interbank transfers must be communicated and arranged through automation by a set time each business day. The failure of any vital communication could cost the bank large sums in interest charged by the central bank. A BAM solution would be programmed to become aware of each message and await confirmation. Failure to receive confirmation within a reasonable amount of time would be grounds for the BAM solution to raise an alarm that would set in motion manual intervention to investigate the cause of the delay and to push the problem toward resolution before it becomes costly.

Another example involves a mobile telecommunications company interested in detecting a situation whereby new customers are not set up promptly and correctly on their network and within the various CRM and billing solutions. Low-level technical events such as messages passing from one application to another over a middleware system, or transactions detected within a database logfile, are processed by a CEP engine. All events relating to an individual customer are correlated in order to detect an anomalous situation whereby an individual customer has not been promptly or correctly provisioned, or set up. An alert can be generated to notify technical operations or to notify business operations, and the failed provisioning step may be restarted automatically.

Complex Event Processing

Event processing is a method of tracking and analyzing (processing) streams of information (data) about things that happen (events), and deriving a conclusion from them. Complex event processing, or CEP, is event processing that combines data from multiple sources to infer events or patterns that suggest more complicated circumstances. The goal of complex event processing is to identify meaningful events (such as opportunities or threats) and respond to them as quickly as possible.

These events may be happening across the various layers of an organization as sales leads, orders or customer service calls. Or, they may be news items, text messages, social media posts, stock market feeds, traffic reports, weather reports, or other kinds of data. An event may also be defined as a “change of state,” when a measurement exceeds a predefined threshold of time, temperature, or other value. Analysts suggest that CEP will give organizations a new way to analyze patterns in real-time and help the business side communicate better with IT and service departments.

The vast amount of information available about events is sometimes referred to as the event cloud.

Conceptual Description

Among thousands of incoming events, a monitoring system may for instance receive the following three from the same source:

- church bells ringing.
- the appearance of a man in a tuxedo with a woman in a flowing white gown.
- rice flying through the air.

From these events the monitoring system may infer a *complex event*: a wedding. CEP as a technique helps discover complex events by analyzing and correlating other events: the bells, the man and woman in wedding attire and the rice flying through the air.

CEP relies on a number of techniques, including:

- Event-pattern detection
- Event abstraction
- Event filtering
- Event aggregation and transformation

- Modeling event hierarchies
- Detecting relationships (such as causality, membership or timing) between events
- Abstracting event-driven processes

Commercial applications of CEP exist in variety of industries and include algorithmic stock-trading, the detection of credit-card fraud, business activity monitoring, and security monitoring.

History

The CEP area has roots in discrete event simulation, the active database area and some programming languages. The activity in the industry was preceded by a wave of research projects in the 1990s. According to the first project that paved the way to a generic CEP language and execution model was the Rapide project in Stanford University, directed by David Luckham. In parallel there have been two other research projects: Infospheres in California Institute of Technology, directed by K. Mani Chandy, and Apama in University of Cambridge directed by John Bates. The commercial products were dependents of the concepts developed in these and some later research projects. Community efforts started in a series of event processing symposiums organized by the Event Processing Technical Society, and later by the ACM DEBS conference series. One of the community efforts was to produce the event processing manifesto

Related Concepts

CEP is used in Operational Intelligence (OI) solutions to provide insight into business operations by running query analysis against live feeds and event data. OI solutions collect real-time data and correlate against historical data to provide insight into and analysis of the current situation. Multiple sources of data can be combined from different organizational silos to provide a common operating picture that uses current information. Wherever real-time insight has the greatest value, OI solutions can be applied to deliver the information needed.

In network management, systems management, application management and service management, people usually refer instead to event correlation. As CEP engines, event correlation engines (*event correlators*) analyze a mass of events, pinpoint the most significant ones, and trigger actions. However, most of them do not produce new inferred events. Instead, they relate high-level events with low-level events.

Inference engines, e.g. rule-based reasoning engines typically produce inferred information in artificial intelligence. However, they do not usually produce new information in the form of complex (i.e., inferred) events.

Example

A more systemic example of CEP involves a car, some sensors and various events and reactions. Imagine that a car has several sensors—one that measures tire pressure, one that measures speed, and one that detects if someone sits on a seat or leaves a seat.

In the first situation, the car is moving and the pressure of one of the tires moves from 45 psi to 41 psi over 15 minutes. As the pressure in the tire is decreasing, a series of events containing the tire pressure is generated. In addition, a series of events containing the speed of the car is generated. The car's Event Processor may detect a situation whereby a loss of tire pressure over a relatively long period of time results in the creation of the "lossOfTirePressure" event. This new event may trigger a reaction process to note the pressure loss into the car's maintenance log, and alert the driver via the car's portal that the tire pressure has reduced.

In the second situation, the car is moving and the pressure of one of the tires drops from 45 psi to 20 psi in 5 seconds. A different situation is detected—perhaps because the loss of pressure occurred over a shorter period of time, or perhaps because the difference in values between each event were larger than a predefined limit. The different situation results in a new event "blowOutTire" being generated. This new event triggers a different reaction process to immediately alert the driver and to initiate onboard computer routines to assist the driver in bringing the car to a stop without losing control through skidding.

In addition, events that represent detected situations can also be combined with other events in order to detect more complex situations. For example, in the final situation the car is moving normally and suffers a blown tire which results in the car leaving the road and striking a tree, and the driver is thrown from the car. A series of different situations are rapidly detected. The combination of "blowOutTire", "zeroSpeed" and "driverLeftSeat" within a very short period of time results in a new situation being detected: "occupantThrownAccident". Even though there is no direct measurement that can determine conclusively that the driver was thrown, or that there was an accident, the combination of events allows the situation to be detected and a new event to be created to signify the detected situation. This is the essence of a complex (or composite) event. It is complex because one cannot directly detect the situation; one has to infer or deduce that the situation has occurred from a combination of other events.

Types

Most CEP solutions and concepts can be classified into two main categories:

- Aggregation-oriented CEP
- Detection-oriented CEP

An aggregation-oriented CEP solution is focused on executing on-line algorithms as a response to event data entering the system. A simple example is to continuously calculate an average based on data in the inbound events.

Detection-oriented CEP is focused on detecting combinations of events called events patterns or situations. A simple example of detecting a situation is to look for a specific sequence of events.

There also exist hybrid approaches.

Integration with Business Process Management

A natural fit for CEP has been with Business Process Management, or BPM. BPM focuses on end-to-end business processes, in order to continuously optimize and align for its operational environment.

However, the optimization of a business does not rely solely upon its individual, end-to-end processes. Seemingly disparate processes can affect each other significantly. Consider this scenario: In the aerospace industry, it is good practice to monitor breakdowns of vehicles to look for trends (determine potential weaknesses in manufacturing processes, material, etc.). Another separate process monitors current operational vehicles' life cycles and decommissions them when appropriate. One use for CEP is to link these separate processes, so that in the case of the initial process (breakdown monitoring) discovering a malfunction based on metal fatigue (a significant event), an action can be created to exploit the second process (life cycle) to issue a recall on vehicles using the same batch of metal discovered as faulty in the initial process.

The integration of CEP and BPM must exist at two levels, both at the business awareness level (users must understand the potential holistic benefits of their individual processes) and also at the technological level (there needs to be a method by which CEP can interact with BPM implementation). For a recent state of the art review on the integration of CEP with BPM, which is frequently labeled as Event-Driven Business Process Management, refer to.

Computation-oriented CEP's role can arguably be seen to overlap with Business Rule technology.

For example, customer service centers are using CEP for click-stream analysis and customer experience management. CEP software can factor real-time information about millions of events (clicks or other interactions) per second into business intelligence and other decision-support applications. These "recommendation applications" help agents provide personalized service based on each customer's experience. The CEP application may collect data about what customers on the phone are currently doing, or how they have recently interacted with the company in other various channels, including in-branch, or on the Web via self-service features, instant messaging and email. The application then analyzes the total customer experience and recommends scripts or next steps that guide the agent on the phone, and hopefully keep the customer happy.

In Financial Services

The financial services industry was an early adopter of CEP technology, using complex event processing to structure and contextualize available data so that it could inform trading behavior, specifically algorithmic trading, by identifying opportunities or threats that indicate traders (or automatic trading systems) should buy or sell. For example, if a trader wants to track stocks that have five up movements followed by four down movements, CEP technology can track such an event. CEP technology can also track drastic rise and fall in number of trades. Algorithmic trading is already a practice in stock trading. It is estimated that around 60% of Equity trading in the United States is by way of algorithmic trades. CEP is expected to continue to help financial institutions improve their algorithms and be more efficient.

Recent improvements in CEP technologies have made it more affordable, helping smaller firms to create trading algorithms of their own and compete with larger firms. CEP has evolved from an emerging technology to an essential platform of many capital markets. The technology's most consistent growth has been in banking, serving fraud detection, online banking, and multichannel marketing initiatives.

Today, a wide variety of financial applications use CEP, including profit, loss, and risk management systems, order and liquidity analysis, quantitative trading and signal generation systems, and others.

Integration with time series databases

A time series database is a software system that is optimized for the handling of data organized by time. Time series are finite or infinite sequences of data items, where each item has an associated timestamp and the sequence of timestamps is non-decreasing. Elements of a time series are often called ticks. The timestamps are not required to be ascending (merely non-decreasing) because in practice the time resolution of some systems such as financial data sources can be quite low (milliseconds, microseconds or even nanoseconds), so consecutive events may carry equal timestamps.

Time series data provides a historical context to the analysis typically associated with complex event processing. This can apply to any vertical industry such as finance and cooperatively with other technologies such as BPM.

Consider the scenario in finance where there is a need to understand historic price volatility to determine statistical thresholds of future price movements. This is helpful for both trade models and transaction cost analysis.

The ideal case for CEP analysis is to view historical time series and real-time streaming data as a single time continuum. What happened yesterday, last week or last month is simply an extension of what is occurring today and what may occur in the future. An example may involve comparing current market volumes to historic volumes, prices and volatility for trade execution logic. Or the need to act upon live market prices may involve comparisons to benchmarks that include sector and index movements, whose intra-day and historic trends gauge volatility and smooth outliers.

Business Process Management

Business process management (BPM) is a field in operations management that focuses on improving corporate performance by managing and optimizing a company's business processes. It can therefore be described as a "process optimization process." It is argued that BPM enables organizations to be more efficient, more effective and more capable of change than a functionally focused, traditional hierarchical management approach. These processes can impact the cost and revenue generation of an organization.

As a policy-making approach, BPM sees processes as important assets of an organization that must be understood, managed, and developed to announce value-added products and services to clients or customers. This approach closely resembles other total quality management or continual improvement process methodologies and BPM proponents also claim that this approach can be supported, or enabled, through technology. As such, many BPM articles and scholars frequently discuss BPM from one of two viewpoints: people and/or technology.

Definitions

BPMInstitute.org defines Business Process Management as:

> the definition, improvement and management of a firm's end-to-end enterprise business processes in order to achieve three outcomes crucial to a performance-based, customer-driven firm: 1) clarity on strategic direction, 2) alignment of the firm's resources, and 3) increased discipline in daily operations.

The Workflow Management Coalition, BPM.com and several other sources have come to agreement on the following definition:

> Business Process Management (BPM) is a discipline involving any combination of modeling, automation, execution, control, measurement and optimization of business activity flows, in support of enterprise goals, spanning systems, employees, customers and partners within and beyond the enterprise boundaries.

The Association Of Business Process Management Professionals defines BPM as:

> Business Process Management (BPM) is a disciplined approach to identify, design, execute, document, measure, monitor, and control both automated and non-automated business processes to achieve consistent, targeted results aligned with an organization's strategic goals. BPM involves the deliberate, collaborative and increasingly technology-aided definition, improvement, innovation, and management of end-to-end business processes that drive business results, create value, and enable an organization to meet its business objectives with more agility. BPM enables an enterprise to align its business processes to its business strategy, leading to effective overall company performance through improvements of specific work activities either within a specific department, across the enterprise, or between organizations.

Gartner defines Business process management (BPM) as:

> "the discipline of managing processes (rather than tasks) as the means for improving business performance outcomes and operational agility. Processes span organizational boundaries, linking together people, information flows, systems and other assets to create and deliver value to customers and constituents."

It is common to confuse BPM with a BPM Suite (BPMS). BPM is a professional discipline done by people, whereas a BPMS is a technological suite of tools designed to help the BPM professionals accomplish their goals. BPM should also not be confused with an application or solution developed to support a particular process. Suites and solutions represent ways of automating business processes, but automation is only one aspect of BPM.

Changes in Business Process Management

The concept of business process may be as traditional as concepts of tasks, department, production, and outputs, arising from job shop scheduling problems in the early 20th Century. The management and improvement approach as of 2010, with formal definitions and technical modeling, has been around since the early 1990s. Note that the term "business process" is sometimes used by IT practitioners as synonymous with the management of middleware processes or with integrating application software tasks.

Although BPM initially focused on the automation of business processes with the use of information technology, it has since been extended to integrate human-driven processes in which human interaction takes place in series or parallel with the use of technology. For example, workflow management systems can assign individual steps requiring deploying human intuition or judgment to relevant humans and other tasks in a workflow to a relevant automated system.

More recent variations such as "human interaction management" are concerned with the interaction between human workers performing a task.

As of 2010 technology has allowed the coupling of BPM with other methodologies, such as Six Sigma. Some BPM tools such as SIPOCs, process flows, RACIs, CTQs and histograms allow users to:

- visualize - functions and processes
- measure - determine the appropriate measure to determine success
- analyze - compare the various simulations to determine an optimal improvement
- improve - select and implement the improvement
- control - deploy this implementation and by use of user-defined dashboards monitor the improvement in real time and feed the performance information back into the simulation model in preparation for the next improvement iteration
- re-engineer - revamp the processes from scratch for better results

This brings with it the benefit of being able to simulate changes to business processes based on real-world data (not just on assumed knowledge). Also, the coupling of BPM to industry methodologies allows users to continually streamline and optimize the process to ensure that it is tuned to its market need.

As of 2012 research on BPM has paid increasing attention to the compliance of business processes. Although a key aspect of business processes is flexibility, as business processes continuously need to adapt to changes in the environment, compliance with business strategy, policies and government regulations should also be ensured. The compliance aspect in BPM is highly important for governmental organizations. As of 2010 BPM approaches in a governmental context largely focus on operational processes and knowledge representation. Although there have been many technical studies on operational business processes in both the public and private sectors, researchers have rarely taken legal compliance activities into account, for instance the legal implementation processes in public-administration bodies.

BPM Life-cycle

Business process management activities can be arbitrarily grouped into categories such as design, modeling, execution, monitoring, and optimization.

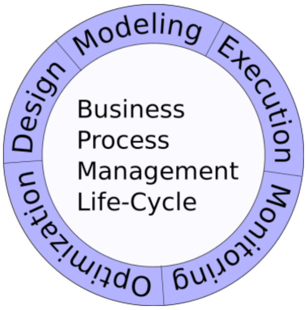

Design

Process design encompasses both the identification of existing processes and the design of "to-be" processes. Areas of focus include representation of the process flow, the factors within it, alerts and notifications, escalations, standard operating procedures, service level agreements, and task hand-over mechanisms.

Whether or not existing processes are considered, the aim of this step is to ensure that a correct and efficient theoretical design is prepared.

The proposed improvement could be in human-to-human, human-to-system or system-to-system workflows, and might target regulatory, market, or competitive challenges faced by the businesses.

The existing process and the design of new process for various applications will have to synchronise and not cause major outage or process interruption.

Modeling

Modeling takes the theoretical design and introduces combinations of variables (e.g., changes in rent or materials costs, which determine how the process might operate under different circumstances).

It may also involve running "what-if analysis"(Conditions-when, if, else) on the processes: *"What if I have 75% of resources to do the same task?" "What if I want to do the same job for 80% of the current cost?"*.

Execution

One of the ways to automate processes is to develop or purchase an application that executes the required steps of the process; however, in practice, these applications rarely execute all the steps of the process accurately or completely. Another approach is to use a combination of software and human intervention; however this approach is more complex, making the documentation process difficult.

As a response to these problems, software has been developed that enables the full business process (as developed in the process design activity) to be defined in a computer language which can be directly executed by the computer. The process models can be run through execution engines that automate the processes directly from the model *(e.g.* calculating a repayment plan for a loan) or, when a step is too complex to automate, Business Process Modeling Notation (BPMN) provides front-end capability for human input. Compared to either of the previous approaches, directly executing a process definition can be more straightforward and therefore easier to improve. However, automating a process definition requires flexible and comprehensive infrastructure, which typically rules out implementing these systems in a legacy IT environment.

Business rules have been used by systems to provide definitions for governing behavior, and a business rule engine can be used to drive process execution and resolution.

Monitoring

Monitoring encompasses the tracking of individual processes, so that information on their state can be easily seen, and statistics on the performance of one or more processes can be provided. An example of this tracking is being able to determine the state of a customer order *(e.g.* order arrived, awaiting delivery, invoice paid) so that problems in its operation can be identified and corrected.

In addition, this information can be used to work with customers and suppliers to improve their connected processes. Examples are the generation of measures on how quickly a customer order is processed or how many orders were processed in the last month. These measures tend to fit into three categories: cycle time, defect rate and productivity.

The degree of monitoring depends on what information the business wants to evaluate and analyze and how business wants it to be monitored, in real-time, near real-time or ad hoc. Here, business activity monitoring (BAM) extends and expands the monitoring tools generally provided by BPMS.

Process mining is a collection of methods and tools related to process monitoring. The aim of process mining is to analyze event logs extracted through process monitoring and to compare them with an *a priori* process model. Process mining allows process analysts to detect discrepancies between the actual process execution and the *a priori* model as well as to analyze bottlenecks.

Optimization

Process optimization includes retrieving process performance information from modeling or monitoring phase; identifying the potential or actual bottlenecks and the potential opportunities for cost savings or other improvements; and then, applying those enhancements in the design of the process. Process mining tools are able to discover critical activities and bottlenecks, creating greater business value.

Re-engineering

When the process becomes too complex or inefficient, and optimization is not fetching the desired output, it is usually recommended by a company steering committee chaired by the president / CEO to re-engineer the entire process cycle. Business process reengineering (BPR) has been used by organizations to attempt to achieve efficiency and productivity at work.

BPM Suites

A market has developed for Enterprise software leveraging the Business Process Management concepts to organize and automate processes. The recent convergence of these software from distinct pieces such as Business rules engine, Business Process Modelling, Business Activity Monitoring and Human Workflow has given birth to integrated Business Process Management Suites. Forrester Research, Inc recognize the BPM suite space through three different lenses:

- human-centric BPM
- integration-centric BPM (Enterprise Service Bus)
- document-centric BPM (Dynamic Case Management)

However, standalone integration-centric and document-centric offerings have matured into separate, standalone markets.

Practice

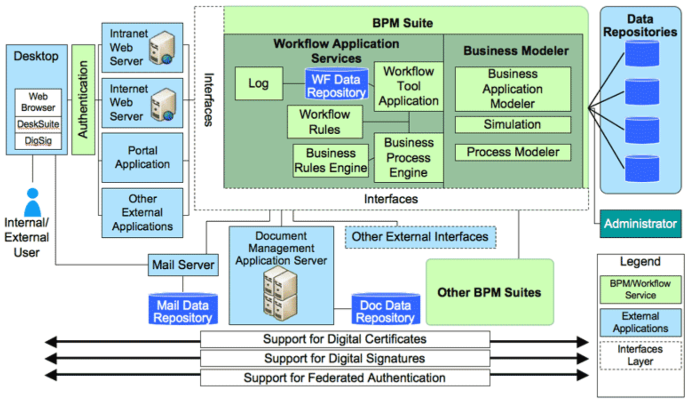

Example of Business Process Management (BPM) Service Pattern: This pattern shows how business process management (BPM) tools can be used to implement business processes through the orchestration of activities between people and systems.

While the steps can be viewed as a cycle, economic or time constraints are likely to limit the process to only a few iterations. This is often the case when an organization uses the approach for short to medium term objectives rather than trying to transform the organizational culture. True iterations are only possible through the collaborative efforts of process participants. In a majority of organizations, complexity will require enabling technology to support the process participants in these daily process management challenges.

To date, many organizations often start a BPM project or program with the objective of optimizing an area that has been identified as an area for improvement.

Currently, the international standards for the task have limited BPM to the application in the IT sector, and ISO/IEC 15944 covers the operational aspects of the business. However, some corporations with the culture of best practices do use standard operating procedures to regulate their operational process. Other standards are currently being worked upon to assist in BPM implementation (BPMN, Enterprise Architecture, Business Motivation Model).

BPM Technology

BPM is now considered a critical component of operational intelligence (OI) solutions to deliver real-time, actionable information. This real-time information can be acted upon in a variety of ways - alerts can be sent or executive decisions can be made using real-time dashboards. OI solutions use real-time information to take automated action based on pre-defined rules so that security measures and or exception management processes can be initiated.

As such, some people view BPM as "the bridge between Information Technology (IT) and Business.". In fact, an argument can be made that this "holistic approach" bridges organizational and technological silos.

There are four critical components of a BPM Suite:

- Process engine — a robust platform for modeling and executing process-based applications, including business rules
- Business analytics — enable managers to identify business issues, trends, and opportunities with reports and dashboards and react accordingly
- Content management — provides a system for storing and securing electronic documents, images, and other files
- Collaboration tools — remove intra- and interdepartmental communication barriers through discussion forums, dynamic workspaces, and message boards

BPM also addresses many of the critical IT issues underpinning these business drivers, including:

- Managing end-to-end, customer-facing processes
- Consolidating data and increasing visibility into and access to associated data and information
- Increasing the flexibility and functionality of current infrastructure and data
- Integrating with existing systems and leveraging service oriented architecture (SOA)
- Establishing a common language for business-IT alignment

Validation of BPMS is another technical issue that vendors and users need to be aware of, if regulatory compliance is mandatory. The validation task could be performed either by an authenticated third party or by the users themselves. Either way, validation documentation will need to be generated. The validation document usually can either be published officially or retained by users.

Cloud Computing BPM

Cloud computing business process management is the use of (BPM) tools that are delivered as software services (SaaS) over a network. Cloud BPM business logic is deployed on an application server and the business data resides in cloud storage.

Market

According to Gartner, 20% of all the "shadow business processes" will be supported by BPM cloud platforms. Gartner refers to all the hidden organizational processes that are supported by IT departments as part of legacy business processes such as Excel spreadsheets, routing of emails using rules, phone calls routing, etc. These can, of course also be replaced by other technologies such as workflow software.

Benefits

The benefits of using cloud BPM services include removing the need and cost of maintaining specialized technical skill sets in-house and reducing distractions from an enterprise's main focus. It offers controlled IT budgeting and enables geographical mobility..

Internet of Things

The emerging Internet of Things poses a significant challenge to control and manage the flow of information through large numbers of devices. To cope with this, a new direction known as BPM Everywhere shows promise as way of blending traditional process techniques, with additional capabilities to automate the handling of all the independent devices.

Metadata

In the 2010s, metadata typically refers to digital forms; however, even traditional card catalogues from the 1960s and 1970s are an example of metadata, as the cards contain information about the books in the library (author, title, subject, etc.).

Metadata is "data [information] that provides information about other data". Three distinct types of metadata exist: structural metadata, descriptive metadata, and administrative metadata.

Structural metadata is data about the containers of data. For instance a "book" contains data, and data about the book is metadata about that container of data.

Descriptive metadata uses individual instances of application data or the data content.

In many countries, the metadata relating to emails, telephone calls, web pages, video traffic, IP connections and cell phone locations are routinely stored by government organizations.

History

Metadata was traditionally used in the card catalogs of libraries until the 1980s, when libraries converted their catalog data to digital databases. In the 2000s, as digital formats are becoming the prevalent way of storing data and information, metadata is also used to describe digital data using metadata standards.

There are different metadata standards for each different discipline (e.g., museum collections, digital audio files, websites, etc.). Describing the contents and context of data or data files increases its usefulness. For example, a web page may include metadata specifying what software language the page is written in (e.g., HTML), what tools were used to create it, what subjects the page is about, and where to find more information about the subject. This metadata can automatically improve the reader's experience and make it easier for users to find the web page online. A CD may include metadata providing information about the musicians, singers and songwriters whose work appears on the disc.

A principal purpose of metadata is to help users find relevant information and discover resources. Metadata also helps to organize electronic resources, provide digital identification, and support the archiving and preservation of resources. Metadata assists users in resource discovery by "allowing resources to be found by relevant criteria, identifying resources, bringing similar resources together, distinguishing dissimilar resources, and giving location information." Metadata of telecommunication activities including Internet traffic is very widely collected by various national governmental organizations. This data is used for the purposes of traffic analysis and can be used for mass surveillance.

Definition

Metadata means "data about data". Although the "meta" prefix means «after» or «beyond», it is used to mean «about» in epistemology. Metadata is defined as the data providing information about one or more aspects of the data; it is used to summarize basic information about data which can make tracking and working with specific data easier. Some examples include:

- Means of creation of the data
- Purpose of the data
- Time and date of creation
- Creator or author of the data
- Location on a computer network where the data was created
- Standards used
- File size

For example, a digital image may include metadata that describes how large the picture is, the color depth, the image resolution, when the image was created, the shutter speed, and other data. A text document's metadata may contain information about how long the document is, who the author is, when the document was written, and a short summary of the document. Metadata within web pages can also contain descriptions of page content, as well as key words linked to the content. These links are often called "Metatags", which were used as the primary factor in determining order for a web search until the late 1990s. The reliance of metatags in web searches was decreased in the late 1990s because of "keyword stuffing". Metatags were being largely misused to trick search engines into thinking some websites had more relevance in the search than they really did.

Metadata can be stored and managed in a database, often called a metadata registry or metadata repository. However, without context and a point of reference, it might be impossible to identify metadata just by looking at it. For example: by itself, a database containing several numbers, all 13 digits long could be the results of calculations or a list of numbers to plug into an equation - without any other context, the numbers themselves can be perceived as the data. But if given the context that this database is a log of a book collection, those 13-digit numbers may now be identified as ISBNs - information that refers to the book, but is not itself the information within the book. The term "metadata" was coined in 1968 by Philip Bagley, in his book "Extension of Programming Language Concepts" where it is clear that he uses the term in the ISO 11179 "traditional" sense, which is "structural metadata" i.e. "data about the containers of data"; rather than the alternate sense "content about individual instances of data content" or metacontent, the type of data usually found in library catalogues. Since then the fields of information management, information science, information technology, librarianship, and GIS have widely adopted the term. In these fields the word *metadata* is defined as "data about data". While this is the generally accepted definition, various disciplines have adopted their own more specific explanation and uses of the term.

Types

While the metadata application is manifold, covering a large variety of fields, there are specialized and well-accepted models to specify types of metadata. Bretherton & Singley (1994) distinguish between two distinct classes: structural/control metadata and guide metadata. *Structural metadata* describes the structure of database objects such as tables, columns, keys and indexes. *Guide metadata* helps humans find specific items and are usually expressed as a set of keywords in a natural language. According to Ralph Kimball metadata can be divided into 2 similar categories: technical metadata and business metadata. *Technical metadata* corresponds to internal metadata, and *business metadata* corresponds to external metadata. Kimball adds a third category, *process metadata*. On the other hand, NISO distinguishes among three types of metadata: descriptive, structural, and administrative.

Descriptive metadata is typically used for discovery and identification, as information to search and locate an object, such as title, author, subjects, keywords, publisher. *Structural metadata* describes how the components of an object are organized. An example of structural metadata would be how pages are ordered to form chapters of a book. Finally, *administrative metadata* gives information to help manage the source. Administrative metadata refers to the technical information, including file type, or when and how the file was created. Two sub-types of administrative metadata are rights management metadata and preservation metadata. *Rights management metadata* explains intellectual property rights, while *preservation metadata* contains information to preserve and save a resource.

Structures

Metadata (metacontent) or, more correctly, the vocabularies used to assemble metadata (metacontent) statements, is typically structured according to a standardized concept using a well-defined metadata scheme, including: metadata standards and metadata models. Tools such as controlled vocabularies, taxonomies, thesauri, data dictionaries, and metadata registries can be used to apply further standardization to the metadata. Structural metadata commonality is also of paramount importance in data model development and in database design.

Syntax

Metadata (metacontent) syntax refers to the rules created to structure the fields or elements of metadata (metacontent). A single metadata scheme may be expressed in a number of different markup or programming languages, each of which requires a different syntax. For example, Dublin Core may be expressed in plain text, HTML, XML, and RDF.

A common example of (guide) metacontent is the bibliographic classification, the subject, the Dewey Decimal class number. There is always an implied statement in any "classification" of some object. To classify an object as, for example, Dewey class number 514 (Topology) (i.e. books having the number 514 on their spine) the implied statement is: "<book><subject heading><514>. This is a subject-predicate-object triple, or more importantly, a class-attribute-value triple. The first two elements of the triple (class, attribute) are pieces of some structural metadata having a defined semantic. The third element is a value, preferably from some controlled vocabulary, some reference (master) data. The combination of the metadata and master data elements results in a statement which is a metacontent statement i.e. "metacontent = metadata + master data". All of these elements can be thought of as "vocabulary". Both metadata and master data are vocabularies which can be assembled into metacontent statements. There are many sources of these vocabularies, both meta and master data: UML, EDIFACT, XSD, Dewey/UDC/LoC, SKOS, ISO-25964, Pantone, Linnaean Binomial Nomenclature, etc. Using controlled vocabularies for the components of metacontent statements, whether for indexing or finding, is endorsed by ISO 25964: "If both the indexer and the searcher are guided to choose the same term for the same concept, then relevant documents will be retrieved." This is particularly relevant when considering search engines of the internet, such as Google. The process indexes pages then matches text strings using its complex algorithm; there is no intelligence or "inferencing" occurring, just the illusion thereof.

Hierarchical, Linear and Planar Schemata

Metadata schemata can be hierarchical in nature where relationships exist between metadata elements and elements are nested so that parent-child relationships exist between the elements. An example of a hierarchical metadata schema is the IEEE LOM schema, in which metadata elements may belong to a parent metadata element. Metadata schemata can also be one-dimensional, or linear, where each element is completely discrete from other elements and classified according to one dimension only. An example of a linear metadata schema is the Dublin Core schema, which is one dimensional. Metadata schemata are often two dimensional, or planar, where each element is completely discrete from other elements but classified according to two orthogonal dimensions.

Hypermapping

In all cases where the metadata schemata exceed the planar depiction, some type of hypermapping is required to enable display and view of metadata according to chosen aspect and to serve special views. Hypermapping frequently applies to layering of geographical and geological information overlays.

Granularity

The degree to which the data or metadata is structured is referred to as its "granularity". "Granularity" refers to how much detail is provided. Metadata with a high granularity allows for deeper, more detailed, and more structured information and enables greater levels of technical manipulation. A lower level of granularity means that metadata can be created for considerably lower costs but will not provide as detailed information. The major impact of granularity is not only on creation and capture, but moreover on maintenance costs. As soon as the metadata structures become outdated, so too is the access to the referred data. Hence granularity must take into account the effort to create the metadata as well as the effort to maintain it.

Standards

International standards apply to metadata. Much work is being accomplished in the national and international standards communities, especially ANSI (American National Standards Institute) and ISO (International Organization for Standardization) to reach consensus on standardizing metadata and registries. The core metadata registry standard is ISO/IEC 11179 Metadata Registries (MDR), the framework for the standard is described in ISO/IEC 11179-1:2004. A new edition of Part 1 is in its final stage for publication in 2015 or early 2016. It has been revised to align with the current edition of Part 3, ISO/IEC 11179-3:2013 which extends the MDR to support registration of Concept Systems. This standard specifies a schema for recording both the meaning and technical structure of the data for unambiguous usage by humans and computers. ISO/IEC 11179 standard refers to metadata as information objects about data, or "data about data". In ISO/IEC 11179 Part-3, the information objects are data about Data Elements, Value Domains, and other reusable semantic and representational information objects that describe the meaning and technical details of a data item. This standard also prescribes the details for a metadata registry, and for registering and administering the information objects within a Metadata Registry. ISO/IEC 11179 Part 3 also has provisions for describing compound structures that are derivations of other data elements, for example through calculations, collections of one or more data elements, or other forms of derived data. While this standard describes itself originally as a "data element" registry, its purpose is to support describing and registering metadata content independently of any particular application, lending the descriptions to being discovered and reused by humans or computers in developing new applications, databases, or for analysis of data collected in accordance with the registered metadata content. This standard has become the general basis for other kinds of metadata registries, reusing and extending the registration and administration portion of the standard.

The Dublin Core metadata terms are a set of vocabulary terms which can be used to describe resources for the purposes of discovery. The original set of 15 classic metadata terms, known as the Dublin Core Metadata Element Set are endorsed in the following standards documents:

- IETF RFC 5013
- ISO Standard 15836-2009
- NISO Standard Z39.85.

Although not a standard, Microformat (also mentioned in the section metadata on the internet below) is a web-based approach to semantic markup which seeks to re-use existing HTML/XHTML tags to convey metadata. Microformat follows XHTML and HTML standards but is not a standard in itself. One advocate of microformats, Tantek Çelik, characterized a problem with alternative approaches:

> “Here’s a new language we want you to learn, and now you need to output these additional files on your server. It’s a hassle. (Microformats) lower the barrier to entry.”

Use

Photographs

Metadata may be written into a digital photo file that will identify who owns it, copyright and contact information, what brand or model of camera created the file, along with exposure information (shutter speed, f-stop, etc.) and descriptive information, such as keywords about the photo, making the file or image searchable on a computer and/or the Internet. Some metadata is created by the camera and some is input by the photographer and/or software after downloading to a computer. Most digital cameras write metadata about model number, shutter speed, etc., and some enable you to edit it; this functionality has been available on most Nikon DSLRs since the Nikon D3, on most new Canon cameras since the Canon EOS 7D, and on most Pentax DSLRs since the Pentax K-3. Metadata can be used to make organizing in post-production easier with the use of key-wording. Filters can be used to analyze a specific set of photographs and create selections on criteria like rating or capture time.

Photographic Metadata Standards are governed by organizations that develop the following standards. They include, but are not limited to:

- IPTC Information Interchange Model IIM (International Press Telecommunications Council),
- IPTC Core Schema for XMP
- XMP – Extensible Metadata Platform (an ISO standard)
- Exif – Exchangeable image file format, Maintained by CIPA (Camera & Imaging Products Association) and published by JEITA (Japan Electronics and Information Technology Industries Association)

- Dublin Core (Dublin Core Metadata Initiative – DCMI)
- PLUS (Picture Licensing Universal System).
- VRA Core (Visual Resource Association)

Telecommunications

Information on the times, origins and destinations of phone calls, electronic messages, instant messages and other modes of telecommunication, as opposed to message content, is another form of metadata. Bulk collection of this call detail record metadata by intelligence agencies has proven controversial after disclosures by Edward Snowden Intelligence agencies such as the NSA are keeping online metadata of millions of internet user for up to a year, regardless of whether or not they are persons of interest to the agency.

Video

Metadata is particularly useful in video, where information about its contents (such as transcripts of conversations and text descriptions of its scenes) is not directly understandable by a computer, but where efficient search of the content is desirable. There are two sources in which video metadata is derived: (1) operational gathered metadata, that is information about the content produced, such as the type of equipment, software, date, and location; (2) human-authored metadata, to improve search engine visibility, discoverability, audience engagement, and providing advertising opportunities to video publishers. In today's society most professional video editing software has access to metadata. Avid's MetaSync and Adobe's Bridge are two prime examples of this.

Web Pages

Web pages often include metadata in the form of meta tags. Description and keywords in meta tags are commonly used to describe the Web page's content. Meta elements also specify page description, key words, authors of the document, and when the document was last modified. Web page metadata helps search engines and users to find the types of web pages they are looking for.

Creation

Metadata can be created either by automated information processing or by manual work. Elementary metadata captured by computers can include information about when an object was created, who created it, when it was last updated, file size, and file extension. In this context an *object* refers to any of the following:

- A physical item such as a book, CD, DVD, a paper map, chair, table, flower pot, etc.
- An electronic file such as a digital image, digital photo, electronic document, program file, database table, etc.

Data virtualization has emerged in the 2000s as the new software technology to complete the virtualization "stack" in the enterprise. Metadata is used in data virtualization servers which are enterprise infrastructure components, alongside database and application servers. Metadata in these servers is saved as persistent repository and describe business objects in various enterprise systems and applications. Structural metadata commonality is also important to support data virtualization.

Statistics and Census Services

Standardization work has had a large impact on efforts to build metadata systems in the statistical community. Several metadata standards are described, and their importance to statistical agencies is discussed. Applications of the standards at the Census Bureau, Environmental Protection Agency, Bureau of Labor Statistics, Statistics Canada, and many others are described. Emphasis is on the impact a metadata registry can have in a statistical agency.

Library and Information Science

Metadata has been used in various ways as a means of cataloging items in libraries in both digital and analog format. Such data helps classify, aggregate, identify, and locate a particular book, DVD, magazine or any object a library might hold in its collection. Until the 1980s, many library catalogues used 3x5 inch cards in file drawers to display a book's title, author, subject matter, and an abbreviated alpha-numeric string (call number) which indicated the physical location of the book within the library's shelves. The Dewey Decimal System employed by libraries for the classification of library materials by subject is an early example of metadata usage. Beginning in the 1980s and 1990s, many libraries replaced these paper file cards with computer databases. These computer databases make it much easier and faster for users to do keyword searches. Another form of older metadata collection is the use by US Census Bureau of what is known as the "Long Form." The Long Form asks questions that are used to create demographic data to find patterns of distribution. Libraries employ metadata in library catalogues, most commonly as part of an Integrated Library Management System. Metadata is obtained by cataloguing resources such as books, periodicals, DVDs, web pages or digital images. This data is stored in the integrated library management system, ILMS, using the MARC metadata standard. The purpose is to direct patrons to the physical or electronic location of items or areas they seek as well as to provide a description of the item/s in question.

More recent and specialized instances of library metadata include the establishment of digital libraries including e-print repositories and digital image libraries. While often based on library principles, the focus on non-librarian use, especially in providing metadata, means they do not follow traditional or common cataloging approaches. Given the custom nature of included materials, metadata fields are often specially created e.g. taxonomic classification fields, location fields, keywords or copyright statement. Standard file information such as file size and format are usually automatically included. Library operation has for decades been a key topic in efforts toward international standardization. Standards for metadata in digital libraries include Dublin Core, METS, MODS, DDI, DOI, URN, PREMIS schema, EML, and OAI-PMH. Leading libraries in the world give hints on their metadata standards strategies.

In Museums

Metadata in a museum context is the information that trained cultural documentation specialists, such as archivists, librarians, museum registrars and curators, create to index, structure, describe, identify, or otherwise specify works of art, architecture, cultural objects and their images. Descriptive metadata is most commonly used in museum contexts for object identification and resource recovery purposes.

Usage

Metadata is developed and applied within collecting institutions and museums in order to:

- Facilitate resource discovery and execute search queries.
- Create digital archives that store information relating to various aspects of museum collections and cultural objects, and serves for archival and managerial purposes.
- Provide public audiences access to cultural objects through publishing digital content online.

Standards

Many museums and cultural heritage centers recognize that given the diversity of art works and cultural objects, no single model or standard suffices to describe and catalogue cultural works. For example, a sculpted Indigenous artifact could be classified as an artwork, an archaeological artifact, or an Indigenous heritage item. The early stages of standardization in archiving, description and cataloging within the museum community began in the late 1990s with the development of standards such as Categories for the Description of Works of Art (CDWA), Spectrum, the Conceptual Reference Model (CIDOC), Cataloging Cultural Objects (CCO) and the CDWA Lite XML schema. These standards use HTML and XML markup languages for machine processing, publication and implementation. The Anglo-American Cataloguing Rules (AACR), originally developed for characterizing books, have also been applied to cultural objects, works of art and architecture. Standards, such as the CCO, are integrated within a Museum's Collection Management System (CMS), a database through which museums are able to manage their collections, acquisitions, loans and conservation. Scholars and professionals in the field note that the "quickly evolving landscape of standards and technologies" create challenges for cultural documentarians, specifically non-technically trained professionals. Most collecting institutions and museums use a relational database to categorize cultural works and their images. Relational databases and metadata work to document and describe the complex relationships amongst cultural objects and multi-faceted works of art, as well as between objects and places, people and artistic movements. Relational database structures are also beneficial within collecting institutions and museums because they allow for archivists to make a clear distinction between cultural objects and their images; an unclear distinction could lead to confusing and inaccurate searches.

Cultural Objects and Art Works

An object's materiality, function and purpose, as well as the size (e.g., measurements, such as height, width, weight), storage requirements (e.g., climate-controlled environment) and focus of the museum and collection, influence the descriptive depth of the data attributed to the object by cultural documentarians. The established institutional cataloging practices, goals and expertise of cultural documentarians and database structure also influence the information ascribed to cultural objects, and the ways in which cultural objects are categorized. Additionally, museums often employ standardized commercial collection management software that prescribes and limits the ways in which archivists can describe artworks and cultural objects. As well, collecting institutions and museums use Controlled Vocabularies to describe cultural objects and artworks in their collections. Getty Vocabularies and the Library of Congress Controlled Vocabularies are reputable within the museum community and are recommended by CCO standards. Museums are encouraged to use controlled vocabularies that are contextual and relevant to their collections and enhance the functionality of their digital information systems. Controlled Vocabularies are beneficial within databases because they provide a high level of consistency, improving resource retrieval. Metadata structures, including controlled vocabularies, reflect the ontologies of the systems from which they were created. Often the processes through which cultural objects are described and categorized through metadata in museums do not reflect the perspectives of the maker communities.

Museums and The Internet

Metadata has been instrumental in the creation of digital information systems and archives within museums, and has made it easier for museums to publish digital content online. This has enabled audiences who might not have had access to cultural objects due to geographic or economic barriers to have access to them. In the 2000s, as more museums have adopted archival standards and created intricate databases, discussions about Linked Data between museum databases have come up in the museum, archival and library science communities. Collection Management Systems (CMS) and Digital Asset Management tools can be local or shared systems. Digital Humanities scholars note many benefits of interoperability between museum databases and collections, while also acknowledging the difficulties achieving such interoperability.

Law

United States of America

Problems involving metadata in litigation in the United States are becoming widespread. Courts have looked at various questions involving metadata, including the discoverability of metadata by parties. Although the Federal Rules of Civil Procedure have only specified rules about electronic documents, subsequent case law has elaborated on the requirement of parties to reveal metadata. In October 2009, the Arizona Supreme Court has ruled that metadata records are public record. Document metadata have proven particularly important in legal environments in which litigation has requested metadata, which can include sensitive information detrimental to a certain party in court. Using metadata removal tools to “clean” or redact documents can mitigate the risks of unwittingly sending sensitive data. This process partially protects law firms from potentially damaging leaking of sensitive data through electronic discovery.

Australia

In Australia the need to strengthen national security has resulted in the introduction of a new metadata storage law. This new law means that both security and policing agencies will be allowed to access up to two years of an individual’s metadata, supposedly to make it easier to stop any terrorist attacks and serious crimes from happening. In the 2000s, the law does not allow access to content of people's messages, phone calls or email and web-browsing history, but these provisions could be changed by the government.

In Healthcare

Australian medical research pioneered the definition of metadata for applications in health care. That approach offers the first recognized attempt to adhere to international standards in medical sciences instead of defining a proprietary standard under the World Health Organization (WHO) umbrella. The medical community yet did not approve the need to follow metadata standards despite research that supported these standards.

Data Warehousing

Data warehouse (DW) is a repository of an organization's electronically stored data. Data warehouses are designed to manage and store the data. Data warehouses differ from business intelligence (BI) systems, because BI systems are designed to use data to create reports and analyze the information, to provide strategic guidance to management. Metadata is an important tool in how data is stored in data warehouses. The purpose of a data warehouse is to house standardized, structured, consistent, integrated, correct, "cleaned" and timely data, extracted from various operational systems in an organization. The extracted data are integrated in the data warehouse environment to provide an enterprise-wide perspective. Data are structured in a way to serve the reporting and analytic requirements. The design of structural metadata commonality using a data modeling method such as entity relationship model diagramming is important in any data warehouse development effort. They detail metadata on each piece of data in the data warehouse. An essential component of a data warehouse/business intelligence system is the metadata and tools to manage and retrieve the metadata. Ralph Kimball describes metadata as the DNA of the data warehouse as metadata defines the elements of the data warehouse and how they work together.

Kimball et al. refers to three main categories of metadata: Technical metadata, business metadata and process metadata. Technical metadata is primarily definitional, while business metadata and process metadata is primarily descriptive. The categories sometimes overlap.

- Technical metadata defines the objects and processes in a DW/BI system, as seen from a technical point of view. The technical metadata includes the system metadata, which defines the data structures such as tables, fields, data types, indexes and partitions in the relational engine, as well as databases, dimensions, measures, and data mining models. Technical metadata defines the data model and the way it is displayed for the users, with the reports, schedules, distribution lists, and user security rights.
- Business metadata is content from the data warehouse described in more user-friendly terms. The business metadata tells you what data you have, where they come from, what they mean and what their relationship is to other data in the data warehouse. Business metadata may also serve as a documentation for the DW/BI system. Users who browse the data warehouse are primarily viewing the business metadata.
- Process metadata is used to describe the results of various operations in the data ware-

house. Within the ETL process, all key data from tasks is logged on execution. This includes start time, end time, CPU seconds used, disk reads, disk writes, and rows processed. When troubleshooting the ETL or query process, this sort of data becomes valuable. Process metadata is the fact measurement when building and using a DW/BI system. Some organizations make a living out of collecting and selling this sort of data to companies - in that case the process metadata becomes the business metadata for the fact and dimension tables. Collecting process metadata is in the interest of business people who can use the data to identify the users of their products, which products they are using, and what level of service they are receiving.

On The Internet

The HTML format used to define web pages allows for the inclusion of a variety of types of metadata, from basic descriptive text, dates and keywords to further advanced metadata schemes such as the Dublin Core, e-GMS, and AGLS standards. Pages can also be geotagged with coordinates. Metadata may be included in the page's header or in a separate file. Microformats allow metadata to be added to on-page data in a way that regular web users do not see, but computers, web crawlers and search engines can readily access. Many search engines are cautious about using metadata in their ranking algorithms due to exploitation of metadata and the practice of search engine optimization, SEO, to improve rankings. See Meta element article for further discussion. This cautious attitude may be justified as people, according to Doctorow, are not executing care and diligence when creating their own metadata and that metadata is part of a competitive environment where the metadata is used to promote the metadata creators own purposes. Studies show that search engines respond to web pages with metadata implementations, and Google has an announcement on its site showing the meta tags that its search engine understands. Enterprise search startup Swiftype recognizes metadata as a relevance signal that webmasters can implement for their website-specific search engine, even releasing their own extension, known as Meta Tags 2.

In Broadcast Industry

In broadcast industry, metadata is linked to audio and video broadcast media to:

- *identify* the media: clip or playlist names, duration, timecode, etc.
- *describe* the content: notes regarding the quality of video content, rating, description (for example, during a sport event, keywords like *goal*, *red card* will be associated to some clips)
- *classify* media: metadata allows to sort the media or to easily and quickly find a video content (a TV news could urgently need some archive content for a subject). For example, the BBC have a large subject classification system, Lonclass, a customized version of the more general-purpose Universal Decimal Classification.

This metadata can be linked to the video media thanks to the video servers. Most major broadcast sport events like FIFA World Cup or the Olympic Games use this metadata to distribute their video content to TV stations through keywords. It is often the host broadcaster who is in charge of organizing metadata through its *International Broadcast Centre* and its video servers. This metadata is recorded with the images and are entered by metadata operators (*loggers*) who associate in live metadata available in *metadata grids* through software (such as Multicam(LSM) or IPDirector used during the FIFA World Cup or Olympic Games).

Geospatial

Metadata that describes geographic objects in electronic storage or format (such as datasets, maps, features, or documents with a geospatial component) has a history dating back to at least 1994 (refer MIT Library page on FGDC Metadata). This class of metadata is described more fully on the geospatial metadata article.

Ecological and Environmental

Ecological and environmental metadata is intended to document the "who, what, when, where, why, and how" of data collection for a particular study. This typically means which organization or institution collected the data, what type of data, which date(s) the data was collected, the rationale for the data collection, and the methodology used for the data collection. Metadata should be generated in a format commonly used by the most relevant science community, such as Darwin Core, Ecological Metadata Language, or Dublin Core. Metadata editing tools exist to facilitate metadata generation (e.g. Metavist, Mercury: Metadata Search System, Morpho). Metadata should describe provenance of the data (where they originated, as well as any transformations the data underwent) and how to give credit for (cite) the data products.

Digital Music

When first released in 1982, Compact Discs only contained a Table Of Contents (TOC) with the number of tracks on the disc and their length in samples. Fourteen years later in 1996, a revision of the CD Red Book standard added CD-Text to carry additional metadata. But CD-Text was not widely adopted. Shortly thereafter, it became common for personal computers to retrieve metadata from external sources (e.g. CDDB, Gracenote) based on the TOC.

Digital audio formats such as digital audio files superseded music formats such as cassette tapes and CDs in the 2000s. Digital audio files could be labelled with more information than could be contained in just the file name. That descriptive information is called the audio tag or audio metadata in general. Computer programs specializing in adding or modifying this information are called tag editors. Metadata can be used to name, describe, catalogue and indicate ownership or copyright for a digital audio file, and its presence makes it much easier to locate a specific audio file within a group, typically through use of a search engine that accesses the metadata. As different digital audio formats were developed, attempts were made to standardize a specific location within the digital files where this information could be stored.

As a result, almost all digital audio formats, including mp3, broadcast wav and AIFF files, have similar standardized locations that can be populated with metadata. The metadata for compressed and uncompressed digital music is often encoded in the ID3 tag. Common editors such as TagLib support MP3, Ogg Vorbis, FLAC, MPC, Speex, WavPack TrueAudio, WAV, AIFF, MP4, and ASF file formats.

Cloud Applications

With the availability of Cloud applications, which include those to add metadata to content, metadata is increasingly available over the Internet.

Administration and Management

Storage

Metadata can be stored either *internally*, in the same file or structure as the data (this is also called *embedded metadata*), or *externally*, in a separate file or field from the described data. A data repository typically stores the metadata *detached* from the data, but can be designed to support embedded metadata approaches. Each option has advantages and disadvantages:

- Internal storage means metadata always travels as part of the data they describe; thus, metadata is always available with the data, and can be manipulated locally. This method creates redundancy (precluding normalization), and does not allow managing all of a system's metadata in one place. It arguably increases consistency, since the metadata is readily changed whenever the data is changed.
- External storage allows collocating metadata for all the contents, for example in a database, for more efficient searching and management. Redundancy can be avoided by normalizing the metadata's organization. In this approach, metadata can be united with the content when information is transferred, for example in Streaming media; or can be referenced (for example, as a web link) from the transferred content. On the down side, the division of the metadata from the data content, especially in standalone files that refer to their source metadata elsewhere, increases the opportunity for misalignments between the two, as changes to either may not be reflected in the other.

Metadata can be stored in either human-readable or binary form. Storing metadata in a human-readable format such as XML can be useful because users can understand and edit it without specialized tools. On the other hand, these formats are rarely optimized for storage capacity, communication time, and processing speed. A binary metadata format enables efficiency in all these respects, but requires special libraries to convert the binary information into human-readable content.

Database Management

Each relational database system has its own mechanisms for storing metadata. Examples of relational-database metadata include:

- Tables of all tables in a database, their names, sizes, and number of rows in each table.
- Tables of columns in each database, what tables they are used in, and the type of data stored in each column.

In database terminology, this set of metadata is referred to as the catalog. The SQL standard specifies a uniform means to access the catalog, called the information schema, but not all databases implement it, even if they implement other aspects of the SQL standard. For an example of database-specific metadata access methods, see Oracle metadata. Programmatic access to metadata is possible using APIs such as JDBC, or SchemaCrawler.

Root Cause Analysis

Root cause analysis (RCA) is a method of problem solving used for identifying the root causes of faults or problems. A factor is considered a root cause if removal thereof from the problem-fault-sequence prevents the final undesirable event from recurring; whereas a causal factor is one that affects an event's outcome, but is not a root cause. Though removing a causal factor can benefit an outcome, it does not prevent its recurrence with certainty.

For example, imagine a fictional segment of students who received poor testing scores. After initial investigation, it was verified that students taking tests in the final period of the school day got lower scores. Further investigation revealed that late in the day, the students lacked ability to focus. Even further investigation revealed that the reason for the lack of focus was hunger. So, the root cause of the poor testing scores was hunger, remedied by moving the testing time to soon after lunch.

As another example, imagine an investigation into a machine that stopped because it overloaded and the fuse blew. Investigation shows that the machine overloaded because it had a bearing that wasn't being sufficiently lubricated. The investigation proceeds further and finds that the automatic lubrication mechanism had a pump which was not pumping sufficiently, hence the lack of lubrication. Investigation of the pump shows that it has a worn shaft. Investigation of why the shaft was worn discovers that there isn't an adequate mechanism to prevent metal scrap getting into the pump. This enabled scrap to get into the pump, and damage it. The root cause of the problem is therefore that metal scrap can contaminate the lubrication system. Fixing this problem ought to prevent the whole sequence of events recurring. Compare this with an investigation that does not find the root cause: replacing the fuse, the bearing, or the lubrication pump will probably allow the machine to go back into operation for a while. But there is a risk that the problem will simply recur, until the root cause is dealt with.

Following the introduction of Kepner–Tregoe analysis—which had limitations in the highly complex arena of rocket design, development and launch—RCA arose in the 1950s as a formal study by the National Aeronautics and Space Administration (NASA) in the United States. New methods of problem analysis developed by NASA included a high level assessment practice called MORT (Management Oversight Risk Tree). MORT differed from RCA by assigning causes to common classes of cause shortcomings that could be summarized into a short list. These included work practice, procedures, management, fatigue, time pressure, along with several others. For example: if an aircraft accident occurred as a result of adverse weather conditions augmented by pressure to leave on time; failure to observe weather precautions could indicate a management or training problem; and lack of appropriate weather concern might indict work practices. Because several measures (methods) may effectively address the root causes of a problem, RCA is an iterative process and a tool of continuous improvement.

RCA is applied to methodically identify and correct the root causes of events, rather than to simply address the symptomatic result. Focusing correction on root causes has the goal of entirely preventing problem recurrence. Conversely, RCFA (Root Cause Failure Analysis) recognizes that complete prevention of recurrence by one corrective action is not always possible.

RCA is typically used as a reactive method of identifying event(s) causes, revealing problems and solving them. Analysis is done *after* an event has occurred. Insights in RCA make it potentially useful as a preemptive method. In that event, RCA can be used to *forecast* or predict probable events even *before* they occur. While one follows the other, RCA is a completely separate process to incident management.

Rather than one sharply defined methodology, RCA comprises many different tools, processes, and philosophies. However, several very-broadly defined approaches or "schools" can be identified by their basic approach or field of origin: safety-based, production-based, assembly-based, process-based, failure-based, and systems-based.

- Safety-based RCA arose from the fields of accident analysis and occupational safety and health.
- Production-based RCA has roots in the field of quality control for industrial manufacturing.
- Process-based RCA, a follow-on to production-based RCA, broadens the scope of RCA to include business processes.
- Failure-based RCA originates in the practice of failure analysis as employed in engineering and maintenance.
- Systems-based RCA has emerged as an amalgam of the preceding schools, incorporating elements from other fields such as change management, risk management and systems analysis.

Despite the different approaches among the various schools of root cause analysis, all share some common principles. Several general processes for performing RCA can also be defined.

General Principles

1. The primary aim of root cause analysis is: to identify the factors that resulted in the nature, the magnitude, the location, and the timing of the harmful outcomes (consequences) of one or more past events; to determine what behaviors, actions, inactions, or conditions need to be changed; to prevent recurrence of similar harmful outcomes; and to identify lessons that may promote the achievement of better consequences. ("Success" is defined as the near-certain prevention of recurrence.)
2. To be effective, root cause analysis must be performed systematically, usually as part of an investigation, with conclusions and root causes that are identified backed up by documented evidence. A team effort is typically required.
3. There may be more than one root cause for an event or a problem, wherefore the difficult part is demonstrating the persistence and sustaining the effort required to determine them.

4. The purpose of identifying all solutions to a problem is to prevent recurrence at lowest cost in the simplest way. If there are alternatives that are equally effective, then the simplest or lowest cost approach is preferred.
5. The root causes identified will depend on the way in which the problem or event is defined. Effective problem statements and event descriptions (as failures, for example) are helpful and usually required to ensure the execution of appropriate analyses.
6. One logical way to trace down root causes is by utilizing hierarchical clustering data-mining solutions (such as graph-theory-based data mining). A root cause is defined in that context as "the conditions that enable one or more causes". Root causes can be deductively sorted out from upper groups of which the groups include a specific cause.
7. To be effective, the analysis should establish a sequence of events or timeline for understanding the relationships between contributory (causal) factors, root cause(s) and the defined problem or event to be prevented.
8. Root cause analysis can help transform a reactive culture (one that reacts to problems) into a forward-looking culture (one that solves problems before they occur or escalate). More importantly, RCA reduces the frequency of problems occurring over time within the environment where the process is used.
9. Root cause analysis as a force for change is a threat to many cultures and environments. Threats to cultures are often met with resistance. Other forms of management support may be required to achieve effectiveness and success with root cause analysis. For example, a "non-punitive" policy toward problem identifiers may be required.

General Process for Performing and Documenting an RCA-based Corrective Action

RCA (in steps 3, 4 and 5) forms the most critical part of successful corrective action, directing the corrective action at the true root cause of the problem. Knowing the root cause is secondary to the goal of prevention, as it is not possible to determine an absolutely effective corrective action for the defined problem without knowing the root cause.

1. Define the problem or describe the event to prevent in the future. Include the qualitative and quantitative attributes (properties) of the undesirable outcomes. Usually this includes specifying the natures, the magnitudes, the locations, and the timing of events. In some cases, "lowering the risks of reoccurrences" may be a reasonable target. For example, "lowering the risks" of future automobile accidents is certainly a more economically attainable goal than "preventing all" future automobile accidents.
2. Gather data and evidence, classifying it along a timeline of events to the final failure or crisis. For every behavior, condition, action and inaction, specify in the "timeline" what should have been done when it differs from what was done.
3. In data mining Hierarchical Clustering models, use the clustering groups instead of classifying: (a) peak the groups that exhibit the specific cause; (b) find their upper-groups; (c) find group characteristics that are consistent; (d) check with experts and validate.

4. Ask "why" and identify the causes associated with each sequential step towards the defined problem or event. "Why" is taken to mean "What were the factors that directly resulted in the effect?"
5. Classify causes into two categories: causal factors that relate to an event in the sequence; and root causes that interrupted that step of the sequence chain when eliminated.
6. Identify all other harmful factors that have equal or better claim to be called "root causes." If there are multiple root causes, which is often the case, reveal those clearly for later optimum selection.
7. Identify corrective action(s) that will, with certainty, prevent recurrence of each harmful effect and related outcomes or factors. Check that each corrective action would, if pre-implemented before the event, have reduced or prevented specific harmful effects.
8. Identify solutions that, when effective and with consensus agreement of the group: prevent recurrence with reasonable certainty; are within the institution's control; meet its goals and objectives; and do not cause or introduce other new, unforeseen problems.
9. Implement the recommended root cause correction(s).
10. Ensure effectiveness by observing the implemented solutions in operation.
11. Identify other possibly useful methodologies for problem solving and problem avoidance.
12. Identify and address the other instances of each harmful outcome and harmful factor.

References

- Luckham, David C. (2012). Event Processing for Business: Organizing the Real-Time Enterprise. Hoboken, New Jersey: John Wiley & Sons, Inc.,. p. 3. ISBN 978-0-470-53485-4.
- National Information Standards Organization; Rebecca Guenther; Jaqueline Radebaugh (2004). Understanding Metadata (PDF). Bethesda, MD: NISO Press. ISBN 1-880124-62-9. Retrieved 2 April 2014.
- Kimball, Ralph (2008). The Data Warehouse Lifecycle Toolkit (Second ed.). New York: Wiley. pp. 10, 115–117, 131–132, 140, 154–155. ISBN 978-0-470-14977-5.
- Wilson, Paul F.; Dell, Larry D.; Anderson, Gaylord F. (1993). Root Cause Analysis: A Tool for Total Quality Management. Milwaukee, Wisconsin: ASQ Quality Press. pp. 8–17. ISBN 0-87389-163-5.
- Taiichi Ohno (1988). Toyota Production System: Beyond Large-Scale Production. Portland, Oregon: Productivity Press. p. 17. ISBN 0-915299-14-3.
- "VRA Core Support Pages". Visual Resource Association Foundation. Visual Resource Association Foundation. Retrieved 27 February 2016.
- Hooland, Seth Van; Verborgh, Ruben (2014). Linked Data for Libraries, Archives and Museums: How to Clean, Link and Publish Your Metadata. London: Facet.
- "ISO 15836:2009 - Information and documentation - The Dublin Core metadata element set". Iso.org. 2009-02-18. Retrieved 2013-08-17.
- Bates, John, John Bates of Progress explains how complex event processing works and how it can simplify the use of algorithms for finding and capturing trading opportunities, Fix Global Trading, retrieved May 14, 2012
- Library of Congress Network Development and MARC Standards Office (2005-09-08). "Library of Congress Washington DC on metadata". Loc.gov. Retrieved 2011-12-23.

Permissions

All chapters in this book are published with permission under the Creative Commons Attribution Share Alike License or equivalent. Every chapter published in this book has been scrutinized by our experts. Their significance has been extensively debated. The topics covered herein carry significant information for a comprehensive understanding. They may even be implemented as practical applications or may be referred to as a beginning point for further studies.

We would like to thank the editorial team for lending their expertise to make the book truly unique. They have played a crucial role in the development of this book. Without their invaluable contributions this book wouldn't have been possible. They have made vital efforts to compile up to date information on the varied aspects of this subject to make this book a valuable addition to the collection of many professionals and students.

This book was conceptualized with the vision of imparting up-to-date and integrated information in this field. To ensure the same, a matchless editorial board was set up. Every individual on the board went through rigorous rounds of assessment to prove their worth. After which they invested a large part of their time researching and compiling the most relevant data for our readers.

The editorial board has been involved in producing this book since its inception. They have spent rigorous hours researching and exploring the diverse topics which have resulted in the successful publishing of this book. They have passed on their knowledge of decades through this book. To expedite this challenging task, the publisher supported the team at every step. A small team of assistant editors was also appointed to further simplify the editing procedure and attain best results for the readers.

Apart from the editorial board, the designing team has also invested a significant amount of their time in understanding the subject and creating the most relevant covers. They scrutinized every image to scout for the most suitable representation of the subject and create an appropriate cover for the book.

The publishing team has been an ardent support to the editorial, designing and production team. Their endless efforts to recruit the best for this project, has resulted in the accomplishment of this book. They are a veteran in the field of academics and their pool of knowledge is as vast as their experience in printing. Their expertise and guidance has proved useful at every step. Their uncompromising quality standards have made this book an exceptional effort. Their encouragement from time to time has been an inspiration for everyone.

The publisher and the editorial board hope that this book will prove to be a valuable piece of knowledge for students, practitioners and scholars across the globe.

Index