Table of Contents for Chapter 4

Part 1. Architecture 101 – 'The Language We Speak'

<<< … >>>

Chapter 4. Application Domains

We build our IT solutions in response to specific needs of businesses or other human activities (alien, or other inhuman activities, are out of the scope of this book).
IT solutions support core business process of the Enterprise, or IT technologies themselves are the core of the business.

There is great variety of areas where we apply our IT Solutions. We call them Application Domains. Specificity of business domains where we apply our IT solutions, and the specificity of IT solutions themselves determine the makeup of Application Domains out there.
IT Solutions for Application Domains in the broad market represent composition or aggregation of tools, techniques, and patterns, components of which may be widely used elsewhere, possibly in other Application Domains.

Taxonomy of Application Domains is settling down but still evolving.
One approach in determining the segmentation of Application Domains for the software is based on industries or business domains – vertical classification.
Another approach looks into generic components of the infrastructure of IT solutions representing ubiquitous building blocks based on some common functionality or capability that Enterprise Solution provides – largely horizontal classification, cutting across various industries.
After all, many industries need Content Management. Not every Enterprise Architecture needs Banking facility.

The more generic, or 'horizontal', is the component or capability, the more potential for re-use, re-purposing, repositioning, and saving.
Business owners of the Enterprise Solution may not always see the longer-term benefits of the good Software Architecture, as long as application functionality fulfils immediate needs of the business.

We've seen some software packages out there that keep the core business going and the customers happy (at least, for a while), but are really ugly under the hood and stretch IT support to the limit. Sooner or later in the life cycle of the application, compromised Software Architecture will surface one way or another – through complex maintenance and enhancements, disappearing market support, diminishing product skills, limited scalability, convoluted integration, "architecture rot", and, finally, increased total cost of ownership.
And, there are examples to the opposite pitfalls, when IT purists build the perfect crystal palace, and lock 'annoying' customers out, so they don't wreak the place. Obviously, unless close collaboration with business marketing and customers is re-established, such project can expect the more immediate problem.

There is only one way in avoiding both these pitfalls on the opposite sides of the spectre – IT Enterprise Solution providers on one hand, and its business owners on another, must establish early and constantly maintain the close working relationships towards the common goal, with common understanding of challenges, and share of responsibilities.

Clear, broadly adopted, and well-understood common vocabulary, or taxonomy for the segmentation and categorisation of the components in the software architecture goes a long way in bridging the gap between Software Architects and business owners of the application.
Both camps need to make an effort to meet halfway. But this book geared more towards Software Architects in the first instance. Understanding of business requirements and the business context is a good start in building the good Software Architecture. After all, who pays for the music?

This chapter outlines some of the major domains and categories in the application software.

Content Management

In response to the needs of modern distributed Enterprise Architectures, Database Management of old evolved through *Document Management* towards the *Content Management*.

Concept of *Content Management* is becoming ubiquitous and generic – every application and enterprise needs one, even if it is not a comprehensive package but set of procedures and guidelines for management of some information assets.
By *Content Management* we mean capture, repository and management of electronic multimedia artefacts that are used throughout the Enterprise Architecture.

Information asset (or *content*) itself may take multitude of shapes and forms (images, video, audio – i.e. various MIME-types), and be complex or compound – like project documentation consisting of many documents.
Two possible approaches in the *Content Management* are based on how we view an information asset – with emphasis on data context or on electronic item itself. Former implies good knowledge about the semantics or the nature of underlying data (and may pull an electronic item apart), latter – treasures an electronic item as an object, entity.
Good *Content Management* system has to find some compromise and satisfy both camps, to the extent that is required by the Enterprise Architecture. This is achieved through building reach metadata – data about data that help us index our information assets and search them.

For example, video library may have an extensive repository of movies. However, we need to know much more than the file name to find the one we want.

Being a data repository and engine in first place, it comes as no surprise that *Content Management* has to deliver CRUD functions for our information assets:

- Recognise and capture an information asset in the repository. This implies creating of some metadata entries as well
- Search, retrieve and render content for our viewing pleasure or for the application needs
- Update content and its metadata
- Delete unwanted content

In addition, *Content Management* in the distributed collaborative environment will provide stringent *Access Control* to the content, as well as *Change Management* and versioning of the content.
Requirements for capacity, bandwidths and low latency for public Internet access are pushing today's' technology and available infrastructure to the limit.

Given variety of capabilities and scale of the Content Management products, and depending on the context and specific needs of your enterprise, finding the *Content Management* product for you – is a big task and an important decision to make.
Really, if you create and maintain a reasonably small public web site with a close-knit team, you will be in different league with large enterprises that require to manage content in the complex and sizeable *Knowledge Management* or e-Commerce portals (Oops, let's take this back – 'portal' is considered by some a dirty word now – too many negative connotations after the dot com bubble burst, and the premature euphoria subsided. We are past the hype-point now. Let's do the real work.)

Knowledge Management

In the most general sense, *Knowledge Management* is about our ability to capture and categorise, digest and maintain, seek and find when needed, our information assets in order to achieve our business goals.
Information assets include not just tangible data, but our experience and internal (often intuitive, not rigorously formalised) or insiders' knowledge of some processes, guidelines and facts in the enterprise. This knowledge includes information on enterprise capability and processes, knowledge of our vendors and suppliers (and their processes and obligations), knowledge of our customers.

There may be three types of knowledge in the organization – individual, group and enterprise. Also, 'that knowledge can be generally classified along the lines of being explicit, embedded, and tacit' [Hummingbird 2001]:

- ***Explicit Knowledge*** – knowledge represented in documents, books, e-mail and databases;
- ***Embedded Knowledge*** – organisational knowledge found in business processes, products and services;
- ***Tacit Knowledge*** – undocumented knowledge that is captured during business processes by knowledge workers.

The overall challenge that many organizations face today is identifying where that knowledge resides and how to leverage it across the enterprise, group and/or individual. The majority of KM initiatives today usually revolve around identifying/discovering, classifying and indexing explicit knowledge in information systems, such as enterprise document management system, and/or business content management system. In many cases KM systems also include access to structured information found in databases [Hummingbird 2001].

Every business success may be traced back to its roots of the successful *Knowledge Management*, and capitalising on this. Implementing of the Knowledge Management crosses many disciplines, technical or organisational, and is equally important concern for every stakeholder in the enterprise.

Knowledge is an asset that is often kept to ourselves. Whole institution of consultancy is based on selling the knowledge and expertise (and, sometimes, just a perception of it) at a profit.
Do not make a mistake about it – the very notion of sharing the knowledge may cause a culture shock and resistance in the enterprise.

In a modern enterprise, knowledge and the 'know-how' is a major differentiating factor for the successful business. For the enterprise to succeed, it has to recognise, treasure, cultivate, nurture and grow its knowledge base, and make it readily and easily accessible in daily needs of the business, in collaborative fashion.
Knowledge Management is becoming an intrinsic part of the Enterprise Architecture.

Information Technology is a great catalyst for the *Knowledge Management* in the enterprise.
IT provides vast storage for repository of data, management of the information assets (including 'behind the scene' integration with applications-consumers of the precious knowledge) and superior to the human (at least in some not overly creative aspects of it) number-crunching capability.

As such, any IT Solutions and Services provider may rightfully claim that they are in the business of *Knowledge Management*, to extent. And this service is always in demand.
Sounds like we belong to the old and safest profession… Do not get carried away – we meant panel-beaters.

Building and implementing the *Knowledge Management* software and systems demand great effort and commitment from all stakeholders.
Given confusing variety of *Knowledge Management* products on the market (and blurred boundaries of the market itself) it is critical to identify core needs of your business and to find the best offering that fits for purpose.

Major players in the Knowledge Management market are evolving from the earlier and better-established markets:

- Data Management and Search;
- Collaboration and Groupware;
- Document management and Content Management;
- Innovative Knowledge Management and Business Intelligence vendors

Knowledge Management is here to stay, but watch the hype and do not expect a silver bullet.

Customer Relationship Management (CRM)

We can be very proud of achievements and advancements in our Enterprise.

However, if paying customer does not share our sentiments, we cannot stay in a dream world much longer.

Customer Relationship Management (CRM) is a front-end of the Enterprise, its front door or its customer-facing side. Obviously, there is no use for front door if there is no house – integration to the Enterprise back-end or 'back office' is of paramount importance.
CRM is directly concerned with servicing, communicating with, satisfying and retaining (and the building up) the Customer Base in a cost-effective manner.

Customer Relationship Management may be referred to as *Customer Interaction Management* (CIM) as well.
Until recently, CIM was almost synonymous to *Call Centre*, implying that customer talks to the human operator to seek help or conduct his or hers business transactions. With technologies like WAP, email, IVR – this is not longer the case. For the same reason, *Call Centres* evolve into *Contact Centres*, with multi-channel and integrated access paths to the Enterprise.

CRM seeks providing customer with access to the Enterprise for the business transactions any time, from any place, any way.
Internet and distributed technologies are only a foundation for the CRM – knock in the customer's door. Some applications rely on premise that 'if you build it, they'll come'. CRM shall overcome this complacency and provide Enterprise Architecture with infrastructure and processes that make a special effort of attracting the large and diverse groups of target users into the Customer Base, and retaining them there.
If Enterprise wants to see a valued customer in the 'front door', it has to provide the tangible value to the customer in the form of convenience, competitive cost structure, peace of mind, diversity and quality of service. Delivery of multiple access channels, security, useability, and high availability goes a long way towards achieving the customer satisfaction and buy-in.

As a customer-facing Enterprise Architecture component, CRM is a primary stakeholder and beneficiary in the quality of *User Interface* design. In order to reach its customers, Enterprise shall know them well and target them in full knowledge of users' tastes, skills, expectations, as well as User Interface standards and legal requirements.

Enterprise Resource Planning (ERP)

Enterprise Resource Planning encompasses two main components of the Enterprise Architecture – Manufacturing and Supply Chain Management.

ERP evolved from *Materials Requirements Planning* (MRP) and, later, *Manufacturing Resource Planning* (MRP II).
MRP determined requirements for materials and scheduled deliveries to keep the production line going non-stop. MRP II has done one better – balanced production capacity and stock of materials by optimising inventory and meeting the delivery dates.

ERP goes further than MRP II in supporting the manufacturing processes – integration of financial, manufacturing and supply chain data and processes; better transaction management and presentment of data; workflow facilities.
Level of the manufacturing process automation is largely determined by how deep ERP penetrated the core business processes.

ERP may be integrated with other common components of the manufacturing enterprise – *Process Control Systems* (PCS) and *Computer Aided Design* (CAD), *Computer Aided Engineering* (CAE), *Computer Aided Manufacturing* (CAM).
CAD systems may be integrated with *Product Data Management* (PDM) that produces documents associated with product development and manufacturing.

PCS manages processes on the factory floor and may include *Programmable Logic Controllers* (PLC), *Supervisory Control and Data Acquisition* (SCADA), *Man-Machine Interface* systems (MMI) and *Material Handling Equipment* (MHE).

Supply Chain Management (SCM)

Supply Chain Management is based on the idea of the industry value chain, meaning that every market or industry has a complete value chain where every company or production step adds value to the product as it goes thru the motions towards its final destination – customer. This is a fundamental concept that promotes integration between ERP systems of various companies in the chain.

Supply Chain Management (SCM) provides IT infrastructure that drives multi-step processes in delivering products and services in the market between multiple providers and consumers of services.

Providers of some services are the consumers of other products and services. Provisioning of complex products and services relies on timely delivery of component products and services that ensure the smooth performance of every player on the market.
Supply Chain Management orchestrates inter-dependent processes and flow of goods in the market so that final composite product is delivered to the satisfied customer, and every link in the supply chain contributed to the fullfillment of the order.

Apart from determining the chain of consumers/providers, and the timing and deliverable of every required service along the way, SCM deals with limited capacity of shared resources and bottlenecks in the chain, with possibly conflicting priorities on every step of the supply process. Resources of consumers/providers in SCM may represent warehouses and distribution centres, production lines, transportation facilities, raw materials, and workforce.
SCM may seek compromises by allowing sub-optimal processes on some steps, provided final deliverable satisfies customer in the end.

Every business transaction has its buyer and seller, consumer and provider. SCM takes a higher-level view by unfolding the whole chain of dependant transactions, and by attempting to manage a bigger picture.
SCM encompasses, and may require close integration of, various IT systems of the participating service providers. Automated management and fullfillment of orders in the complex supply chain necessitates integration of ERP, warehouse, shipment, purchasing, ordering, payment, and other aspects of services provisioning.

SCM may perform two distinct functions – *Supply Chain Planning* (SCP) and *Supply Chain Execution* (SCE).
SCP helps in analysing the *what-if* scenarios before actual commitment of resources and contractual obligations.
SCE kicks in, monitors and helps to control the complex multi-step process of service delivery by conducting the whole orchestra.

Business Intelligence (BI)

Implementation of *Business Intelligence* (BI) framework in the Enterprise is increasingly becoming a necessity and a strong differentiating factor in the IT-intensive environment.

Business Intelligence is about the monitoring and controlling business processes in the enterprise towards achieving business goals. In order to know status of the business and to make well-informed decisions, we need a comprehensive and up-to-date information model of the enterprise.
Business Intelligence framework uses *Data Warehouse* as its foundation.

We cannot manage what we do not understand, or do not know well enough. *Management Reporting* is the heart of the *Business Intelligence*.

In general, *Business Intelligence* frameworks may be categorized along two main dimensions – depth of the captured information and how deep BI tools entrenched into the business processes for faster response to changing business situation.

As such, BI frameworks range from basic *Management Reporting* for further analysis, to automated solutions - when BI framework closes the loop from monitoring to actual management and feedback to the business.

Enterprise Information Portals (EIP)

Enterprise Information Portals (EIP) are the products of convergence between *Knowledge Management* and *Business Intelligence*.

In order to implement a dedicated information delivery function, EIP looks into consolidating the following Enterprise Architecture features:

- Multi-channel access to the conventional 'back office' information repositories, including web-based browser-based access
- Advanced search facility that enables EIP to find, categorise and rank information across multiple information sources
- Document Management and Knowledge Management for collating, authoring, and publishing the valuable content, re-purposing the content for various access channels and user groups
- Collaboration facilities that enable EIP users to communicate and share knowledge, as well as create and manage virtual communities of interest through various forms of collaboration – chat, notice or bulletin boards, forums or threaded discussions (moderated, or not moderated), notifications and alerts, multi-media instant messaging, e-mail
- Security and Access Control, including the Identity Management of the EIP users and provisioning of access rights and privileges to them, in order to enable security and privacy in the information delivery
- Personalisation facility attempts to adapt the user experience so that it is optimised for the user's profile and business goals of the EIP. EIP achieves Personalisation through capturing and tracking user's profile, preferences and behaviours (either explicitly by asking the user, or implicitly by observing the user's activities in EIP and building the profile on the fly, or both)

Data Warehouse

Data Warehouse is yet another incarnation of Data Management, specifically focused on the business decision support.

Data Warehouse is expected to provide a significant competitive advantage to the Enterprise. However, these benefits come with the risks – failure rates have been estimated at as high as up to 70% on projects that can cost millions, and take years to implement.

Data Warehousing is associated with *Datamarts*, *Extract-Transform-Load* processes (ETL), *Data Cleansing* and *Data Mining*, possibly across datastores on various platforms, and across geographically distributed databases.

Datamart represents a somewhat local or departmental database, as opposed to the centralised enterprise-wide *Data Warehouse*. Datamarts are easier to implement and faster to reap rewards of ROI. This seems to be a good idea at times, as long as Enterprise Architect realises what may be sacrificed from the holistic vision for the Enterprise. And, common sacrificial lambs are - metadata for the Enterprise information assets, ease of integration of various enterprise applications, possible duplication of data, and corruption of data integrity in the process.
In order to mitigate risks of making an information island out of datamart, Enterprise Architect has to watch consistency between datamarts and ability for integration of datamarts and applications based on them.

As with *Data Dictionaries* in the databases, metadata (data about data) is the cornerstone of the *Data Warehouse*.

Datamarts often employ technique of dimensional modelling when data categorised along several dimensions or attributes, like product, region, time etc. Data model then can be visualised as a *multi-dimensional cube*. Accordingly, we talk about *Multi-Dimensional Databases* (MDDB).
Note that support of such a nice logical schema or model of the multi-dimensional cube comes at a cost. Different MDDBs may handle physical storage of the cube differently. However, any DBA imagines at ones that you will have to build and maintain indices for every dimension, at least. And, given large volumes of data and possible sparse allocation in the range of key values, this is not a piece of cake.

Term *Data Warehouse*, being quite intuitive, is often abused. In many cases it would be more correct to call data repository for the large business function an *Operational Data Store* – an intermediate data store between raw operational data sources and the *Data Warehouse*.

e-Business and Web Services – B2B, B2C, x2z

Internet is a powerful vehicle for building distributed virtual marketplaces that bring together customers and goods or services providers, consumers and suppliers, partners and businesses, individuals and communities of interest, federated applications and portals.

Depending on the type and makeup of players in e-Business collaboration, we adopted convenient monikers for tipical business scenarious – Business-To-Business (B2B), Business-To-Consumer (B2C), Consumer-To-Consumer (C2C) etc.
In essence, this way we describe the most common patterns in e-Business activities, defined by the criteria of who are the participants in the business transaction.

Note that this definition may describe only one leg in the multi-step business transaction. Other steps along the way in the complex business transaction may have a different *'x2z'* pattern.
For instance, customer may inquire the weather information from the news website ('B2C' scenario). In turn, news website, through some B2B channels, gets hold of the weather information from some weather portal. Furthermore, this weather portal may not be the initial source of all required information… So, the 'B2C' badge does not really stick, and may not satisfactorily describe the nature of the whole e-Business.

In the course of conducting e-Business on the Web, every business transaction has an identifiable client and server. Client makes the use of some service from the server on the Web. Following the intuition, we could say that server in the business transaction provides the client with Web Service. Hold your horses…
With the advent of the XML Web Services (described later in a book), term *Web Services* acquired a particular meaning – kind of specific open-standard, multi-platform, web-based middleware that enables collaboration of heterogeneous systems using XML as a *lingua franka*.
Application function can be developed, published, located, and executed remotely on the Web using XML Web Services on .NET, Java, or any other platform, or combination of any compliant platforms above.

Process Automation using Workflow Processes

Business Process Automation, as seen beyond the scope of the single program, widely relies on the Workflow Management solutions.
Workflows can capture and tie together seemingly disparate or loosely coupled processes, resources, and participants in the Enterprise.

Workflow allows separation of the process definition from the execution of the process, as well as the definition of the whole of the process from the individual steps it consists of.

Workflow standardisation efforts are represented by two main contributions – Workflow Management Coalition (WfMC) Model, and the OMG Workflow Facility.
WfMC Model is early, decidedly non-object-oriented standard, but provides groundwork in defining the workflow components and interfaces.

OMG Workflow Facility acknowledges the fundamentals of the WfMC Reference Model, and provides workflow definitions in terms of IDL and ORB middleware.

Workflow represents some business process. Process may consist of inter-dependant Tasks or Activities. Every Task will have starting and ending point. Process Definition describes pre-conditions for the start of the task, as well as post-conditions upon the exit from the task. Task/Activity (and the whole Process for that matter) may begin its execution when all pre-conditions for this Task are satisfied.
Structure of the Process in the Process Definition is being captured as a network graph of Tasks, where preceding Tasks may trigger some pre-conditions for the following Tasks. Seemingly trivial Task in one Process may be a complex Process on its own – Process Definition is recursive, and may become very complex.
Note - the workflow state after the completion of some preceding Task may trigger pre-conditions in more than one of following Tasks.

In accordance with the nature of business transactions, some business workflows may be fast and furious like a single program. Or, workflow may drag for days and months, if workflow Process happened to be an application for the Credit Card, or lengthy litigation process.

Task or Activity in the workflow can be a whole another workflow Process (*a la* workflow sub-routine), an application program, database process, or even some manual process that may require human intervention (eg. approval signature on the document). Depending on conditions in the workflow during the execution, not every task may be executed in every instance of the workflow. For example, litigation process may have defined the Task of printing and mailing the Eviction Notice. This Task may never be executed in most of workflows, if matter is settled.

Workflow Engine controls the workflow execution. As long as workflow Process Definition is completed and enabled, workflow Process can be instantiated and executed.

Workflow Engine becomes a kind of Transaction Processing Monitor for workflow Processes and Tasks.
Tasks themselves may be external to the Workflow Engine processes, using some external resources. Only requirement mandated by the Workflow Engine is to ensure the Task sends the notification back to Workflow Engine when it is completed, and how – otherwise Workflow Engine cannot trigger following Tasks, and workflow Process may hang.

Workflow Process is a business transaction. Due to often complex and heterogeneous nature of workflow Processes, Workflow Engine may have little control over workflow Tasks and resources. Workflow Engines (as opposed to conventional TPMs and DBMSs) have limited capabilities in enforcing transaction integrity and rollback of workflow Processes.
Workflow Architects must themselves design Process Definitions so that transactional integrity and ACID properties of workflows will not be compromised.

Games and Info-tainment

IT penetrated whole fabric of human society. Computer usage for the ‘serious’ business is the most obvious application of IT.
However, the same ‘serious’ business-people outside business hours (and the hordes of rest of us) represent the vast market of consumers of IT solutions for leisure, but not necessarily for free. And that is ‘serious’ business.
Games and Internet collaboration capability are becoming a major selling point of IT to mass market. Also, info-tainment is becoming a major catalyst for building and delivery of the broadband content to consumer.

Education

We have seen already that IT helps us work hard and play hard. Also, IT helps us learn hard.

Education is a serious and complex enterprise too. From library databases to applications that manage all aspects of the learning processes – preparing, storing and presenting courses; managing student progress and assessments; managing schedules; accounting; consolidating processes across remote campuses – IT feels right at home at School.

<<< … >>>