Motivation: Operational Feasibility

The picture below was the starting point to guide seven Proofs of Concept conducted over 10 months to answer the question, "Are Clouds ready for Enterprise prime time?" Our answer is a qualified "yes" depending on the readiness of the consuming organization. After all, Cloud technologies are old wine in new bottles.

The Open Narrative

Evaluate
Multi-Tenant, Multi-Landlord Platforms
for
Fitness for Business Purpose Environments
IT Operational Sourcing

How Secure is it?

The Questions!

How Hard is it?

How Easy is it to Manage?

How Safe is it?

How is it Monitored?

Figure 1: Questions for Cloud IT Operational Sourcing

The Case for Multi-Landlord Focus: Competitive Pricing

The heavy hype around clouds is its relatively low-cost, multi-tenant aspects. The most troubling question is, however, the security and monitoring models within both external and internal manifestations of clouds.

Moreover, cloud/data center providers (Amazon Web Services, Rackspace, GoGrid, Google App Engine, Microsoft Azure) worry mostly about accommodating multiple consumers (multi-tenant), each shielded from the others within their properties, both real and virtual. Not surprising, providers desire some level of stickiness (read lock-in) of clients.

It was decided immediately, that a dynamically competitive anything-as-a-Service model is most advantageous for consumers. This allows ready switching of providers for most advantageous resource pricing in addition to the elastic quality of resource provisioning.

Thus, multi-landlord was added to the concept of multi-tenant properties. So the question of what are the design patterns with respect to Security and Monitoring issues was asked as illustrated by the classis ER diagram below:

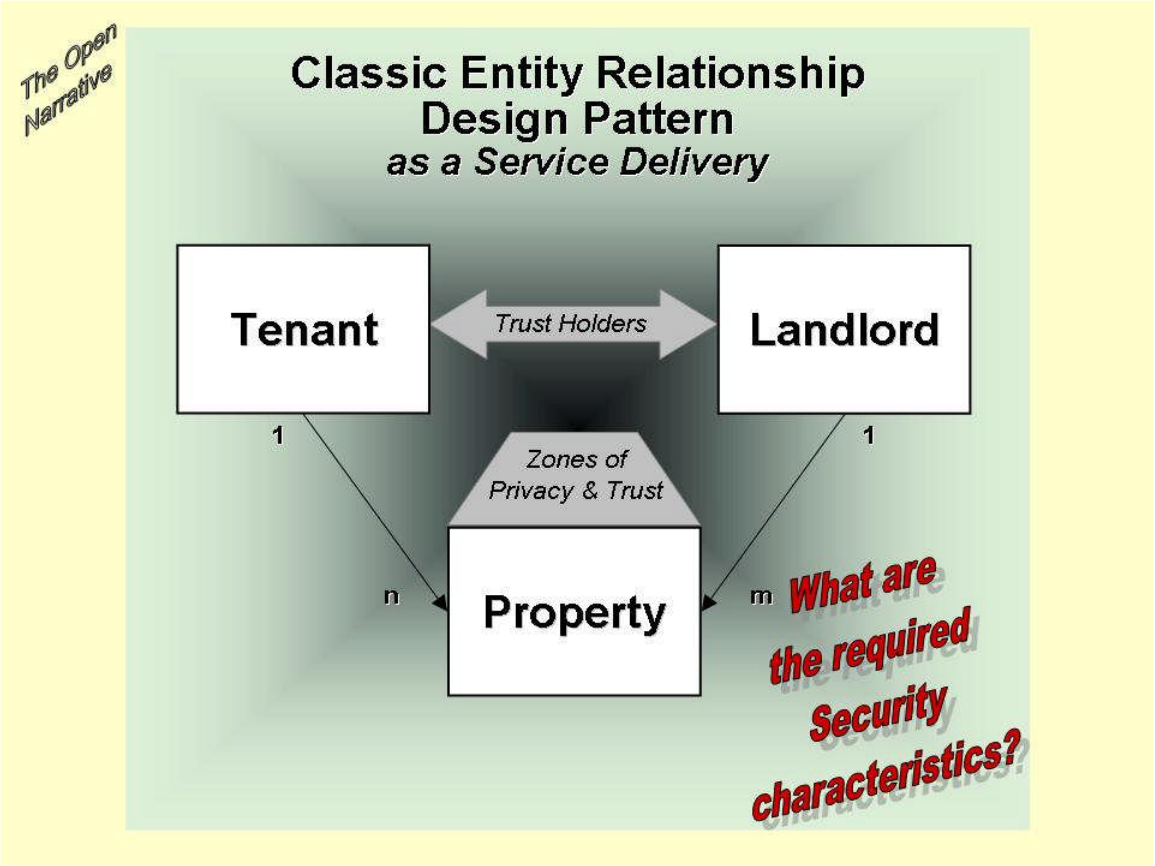

Figure 2: Central Enterprise View of Cloud as a Service Design Pattern

"As a Service" (*aaS) delivery is the operational model embraced by cloud providers. Otherwise, we revert to the old ways of managed services or facilities management. AaS is, for the most part, a recycling and update to the time-sharing models so popular in the '60s and '70s before the commercial advent of workstations and PCs in the '80s.

For security, the above Figure 2 captures the multi-tenant/multi-landlord relationship with respect to the locales of operational facilities, namely, the Zones of Privacy and Trust. The main issues are around engineering the establishment of Trust Domains amongst Trust Holders, Tenants and Landlords, within the Zones of Privacy & Trust.

We turn our spotlight now all the way up the stack on the Data Center and optimizing its provision of fitness for business purpose environments.

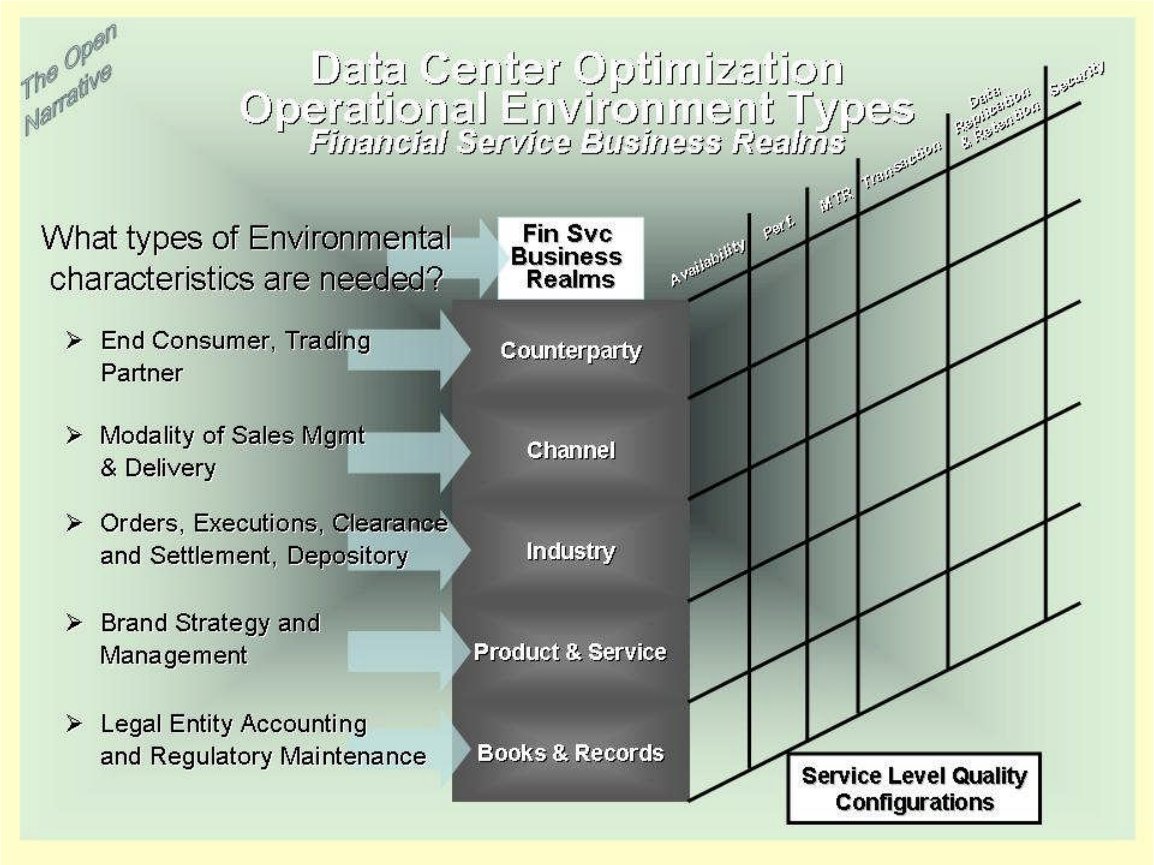
The Open Narrative
Data Center Optimization
Operational Environment Types
Financial Service Business Realms
What types of Environmental characteristics are needed?
➢ End Consumer, Trading Partner
➢ Modality of Sales Mgmt & Delivery
➢ Orders, Executions, Clearance and Settlement, Depository
➢ Brand Strategy and Management
➢ Legal Entity Accounting and Regulatory Maintenance
Fin Svc Business Realms
Counterparty
Channel
Industry
Product & Service
Books & Records
Availability
Perf.
MTR
Transaction
Data Replication & Retention
Security
Service Level Quality Configurations

Figure 3: Financial Services Business Domain Operational Environment Types

Financial Services Business Realms

What are the characteristics of Financial Services Operational Environments to model, tool and experience?

The Financial Services Business Domain Operational Environment Types above in Figure 3 offer a business alignment for Cloud-based Environments of the various Realms of Financial Services.

The Realms represent the demand side of the equation for Cloud supply.

The Environment configurations are around Service Level Quality Requirement aspects such as availability, performance, mean time to recovery, transactional guarantee (once and only once, at least once, best efforts), data replication and retention and, of course, security.

The cost of provisioning Environments is determined, in large part, by the service level quality configuration required to the fitness of an intended business service.

Much of optimizing provisioning costs is in reducing over-provisioned service levels.

Briefly, the Realms are

- **Counterparty**

  Counterparty includes the end consumer, agents for end consumers and trading partners.

  Typical Service Level Requirement:
  Highly Available during market hours, Operate 7X24 for Information Request and Order Entry

- **Channel**

  Channel covers the various modalities of sales management, delivery of offers and fulfilled products and services.

  Typical Service Level Requirement:
  Highly Available during market hours, Open 7X24 for Information Request and Order Entry, Recoverable Sessions

- **Industry**

  Industry is for facing markets with orders, executions, clearance, settlement, and depository.

  Typical Service Level Requirement:
  Highly Available during and after market hours, guaranteed delivery of messages, assured audit ability

- **Products & Services**

  Products and Services are for fulfillment and creating and managing brand strategy and the concomitant offer creations.

  Typical Service Level Requirement:
  Highly Available, Open 7X24 for Information Request, Open Business hours for Service/Product Fulfillment, Non-repudiated transactions

- **Books & Records**

  Books and Records cover legal entity accounting and regulatory oversight.

Typical Service Level Requirement:
Highly Available during quarterly/year-end close, Open business hours for Information Request, Recoverable Sessions

Cloud Models

There are three model types for Clouds:

- **Capability Delivery**

  This model type refers to the level of delivering a particular capability "as a Service."

- **Deployment**

  This model type covers the intended usage patterns of domains of resources and the degree of "openness."

- **Business Purposes**

  This other cloud model type addresses the Enterprise Business Architecture concerns' different classes of collected business functions in the supply chain of products and services, end to end, from User/Client/Customer/Partner/Employee to definition and satisfaction of requests for products and services of the Enterprise all the way to recording and accounting for The Business.

Capability Delivery Models (*aaS)

Capabilities fall into four classes of delivery "as a Service:"

- **Infrastructure**

  This Delivery Model refers to general Information Technology components such as processor, storage, and network as well as system software such as operating systems, database systems, web/application servers.

- **Platform**

  This Delivery Model covers the design centers used to create, operate and maintain applications or business infrastructure. Think of this model as a capability to create end-user toolkit environments for specific business purposes such as CRM, Portfolio Analytics, Collaboration or Social Networking. Platform as a Service is a development abstraction of Software as a Service.

- **Software**

  This Delivery Model provides business functionality for a specific business capability. In a sense, it is an instance of a system created within a Platform as a Service. The most salient example is Salesforce.com and its abstraction as a Platform in Force.com. Interestingly, Force.com was developed after Salesforce.com and does not restrict functionality to Customer Relationship Management.

- **Environment**

  This Delivery Model speaks to fitness for business purposes service level requirements that affect provisioning of business functionality. It abstracts Platform as a Service in the dimensions of Service Level Configuration and Software as a Service as content functionality.

A further note on Environment as a Service

There is current acceptance of the first three—Infrastructure, Platform, and Software— as the major Delivery Models of interest. These three represent the supply side of as-a-Service delivery. Here, we insert and assert Environment as a Service (EaaS) as the demand side, representing the ultimate alignment of technology to business

We further maintain that this is the proper level at which business services will be commoditized (i.e., lowest cost) and delivered as utilities.

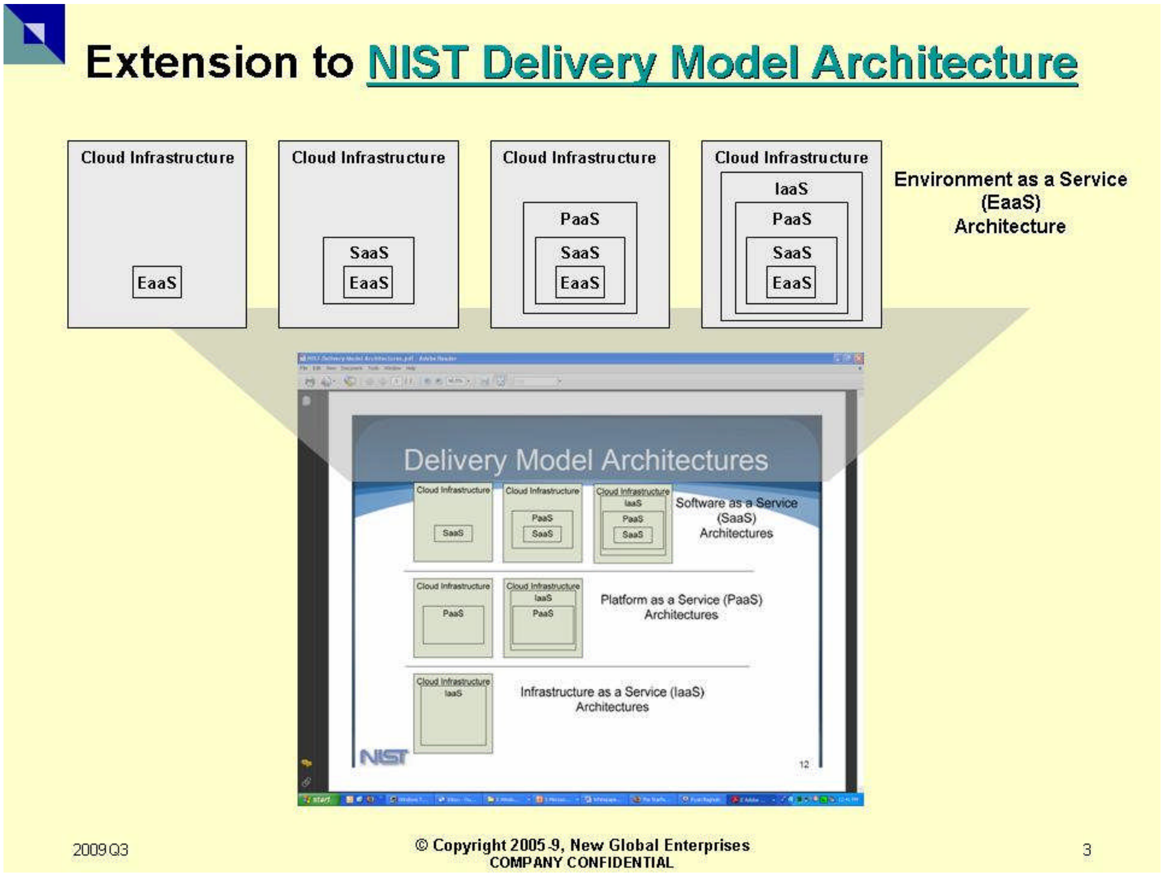

Extension to NIST Delivery Model Architecture

2009Q3

© Copyright 2005-9, New Global Enterprises
COMPANY CONFIDENTIAL

3

Figure 4: Extension to the NIST Delivery Model Architecture

The above diagram, Figure 4, depicts how we extend the Mell-Grance Delivery Model Architecture from page 12 of their seminal NIST thought piece (http://tinyurl.com/lhwlth).

Notice that in this stack of abstractions, as one descends into the innermost box of the rightmost box, one gets more and more recognizable and specific business value. This is worth exploring more in depth, but our purpose here is to explicate the use of clouds to render business production quality operations.

Deployment Models

Deploying clouds is where the proverbial rubber meets the road in operations. In the current debate there are differing views which define what is of interest in modeling cloud deployments. This is vaguely reminiscent of debates over Open Systems 20 years ago. As with Open, we shall eventually converge on a workable meaning. Like the Open meme, Cloud is a persistent meme for IT and shall remain so for the next generation of IT operatives.

The Mell-Grance NIST presentation offers four distinct models from which to choose. Not wishing to recapitulate all the debate and the Mell-Grance treatment, we summarize our understanding of what these deployments are:

- Private Clouds

  While some argue that Private Cloud is an oxymoron; this is where most Enterprises are starting. Namely, Private Cloud represents the next generation of Data Center Optimization from consolidation and virtualization. It is about the internal systems, inside the enterprise firewall, thus conflicts with the concept de-perimeterization that is generally thought as a key characteristic of cloud.

- Community Clouds

  Community is used to cover the social networking capabilities and sites like twitter, facebook or myspace and yahoo mail, hotmail or gmail.

- Public Clouds

  Public refers to the emerging list of exemplary providers like Amazon Web Services, Rackspace or GoGrid delivering IaaS, and, those delivering SaaS/PaaS capabilities like Salesforce.com/Force.com.

- Hybrid Clouds

  Hybrid is a mixture of the above models.

For simplicity, it is argued that Hybrid is a single model used with degrees of access restriction or application focus to define the particular flavor of the above models.

In short, there is one cloud model with variations on the theme to fit the particular circumstance. The advantage of this is (1) simplicity and (2) ease of defining and implementing security policy infrastructure, the focus in this paper.

The position here is there are deployments with configuration specializations that provide the Private, Community, Public or Hybrid behaviors. Behavior is the key, not structure per se.

Business Purpose Models

As noted previously, business value is the prime objective. Intellectual elegance is a nice-to-have. It is business value that determines the got-to-have.

In the previous discussion on deployment models, there is only one needed— the Hybrid Model with variations in configured behaviors. The deployment of instances of clouds then involves a case of federation of Lines of Business operations with Enterprise corporate management control.

That being said, the End2End Business Architecture diagram below appears to be very dense, indeed. As a single Big Picture, it has the merit of driving the discussion of clouds in business value terms that the business can understand and, thus, be willing to fund as they see the value.

Notice well that both delivery and deployment are transparent to this approach. Delivery and deployment are security, monitoring and cost delineations that help make or break a business case.

Consequently, this is business alignment of technology, not vice versa.

With the idea that there is one cloud with a multitude of micro-perimeters, i.e. with lots of nooks and crannies, the discussion below is how to model The Business that operates within The Cloud.

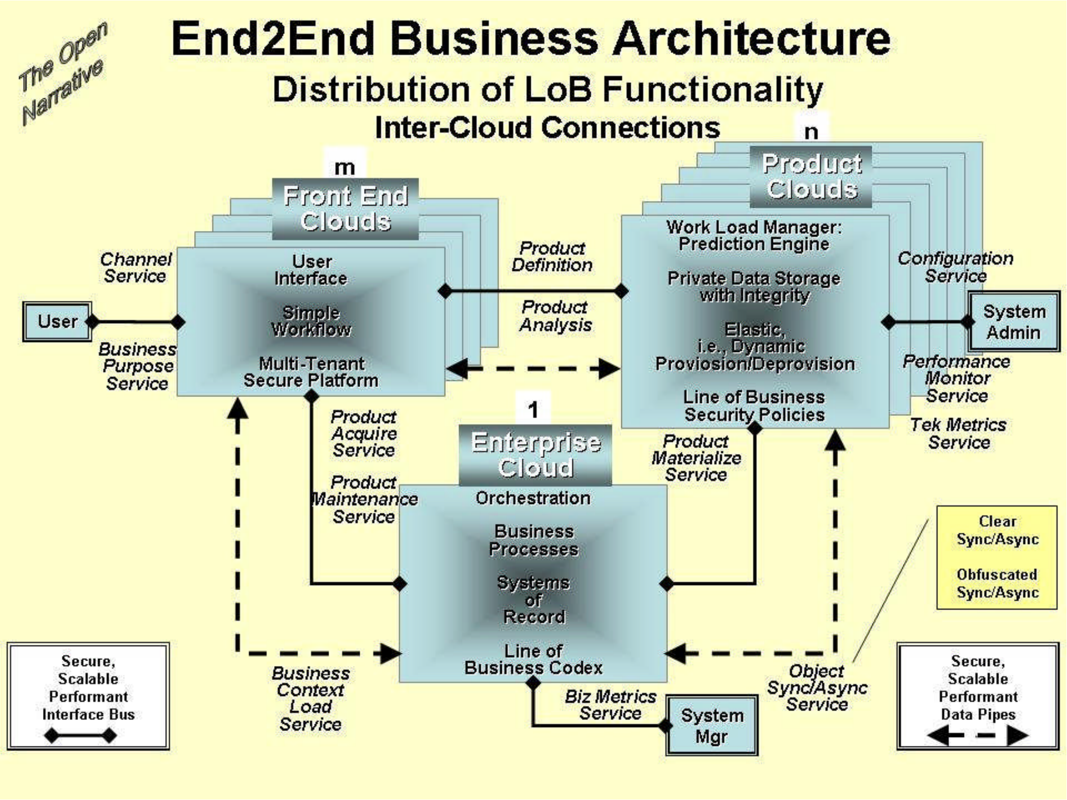

The Open Narrative
End2End Business Architecture
Distribution of LoB Functionality
Inter-Cloud Connections
m
Front End Clouds
User Interface
Simple Workflow
Multi-Tenant Secure Platform
n
Product Clouds
Work Load Manager: Prediction Engine
Private Data Storage with Integrity
Elastic, i.e., Dynamic Proviosion/Deprovision
Line of Business Security Policies
1
Enterprise Cloud
Orchestration
Business Processes
Systems of Record
Line of Business Codex
User
Channel Service
Business Purpose Service
Product Definition
Product Analysis
Configuration Service
System Admin
Performance Monitor Service
Tek Metrics Service
Product Acquire Service
Product Maintenance Service
Product Materialize Service
Business Context Load Service
Biz Metrics Service
System Mgr
Object Sync/Async Service
Clear Sync/Async
Obfuscated Sync/Async
Secure, Scalable Performant Interface Bus
Secure, Scalable Performant Data Pipes

Figure 5: End to End Business Architecture

Types of Business Aligned Clouds

Figure 5 provides an overlay of the three types of Business Oriented Clouds that cover end-to-end business value delivery.

These are enumerated and briefly described below.

Enterprise Clouds

As a definitional entity, Enterprise is where the Books and Records are, the Inner Sanctum—the central control of Orchestration. There IS only one per entity.

Offer Clouds

Profit/non-profit, private/public, Enterprises offer many products and services. In this typology, these are called Offers. Note well that information and influence are construed as resource and service, respectively. This is where multi-landlord drives efficient provisioning of production processes.

Front End Clouds

Enterprises deliver products and service through a variety of channels. This typology envisions clouds for each Front End delivery modality, electronic, human, hybrid. This is where multi-tenancy rules supreme when considering Software of Platform as a Service.

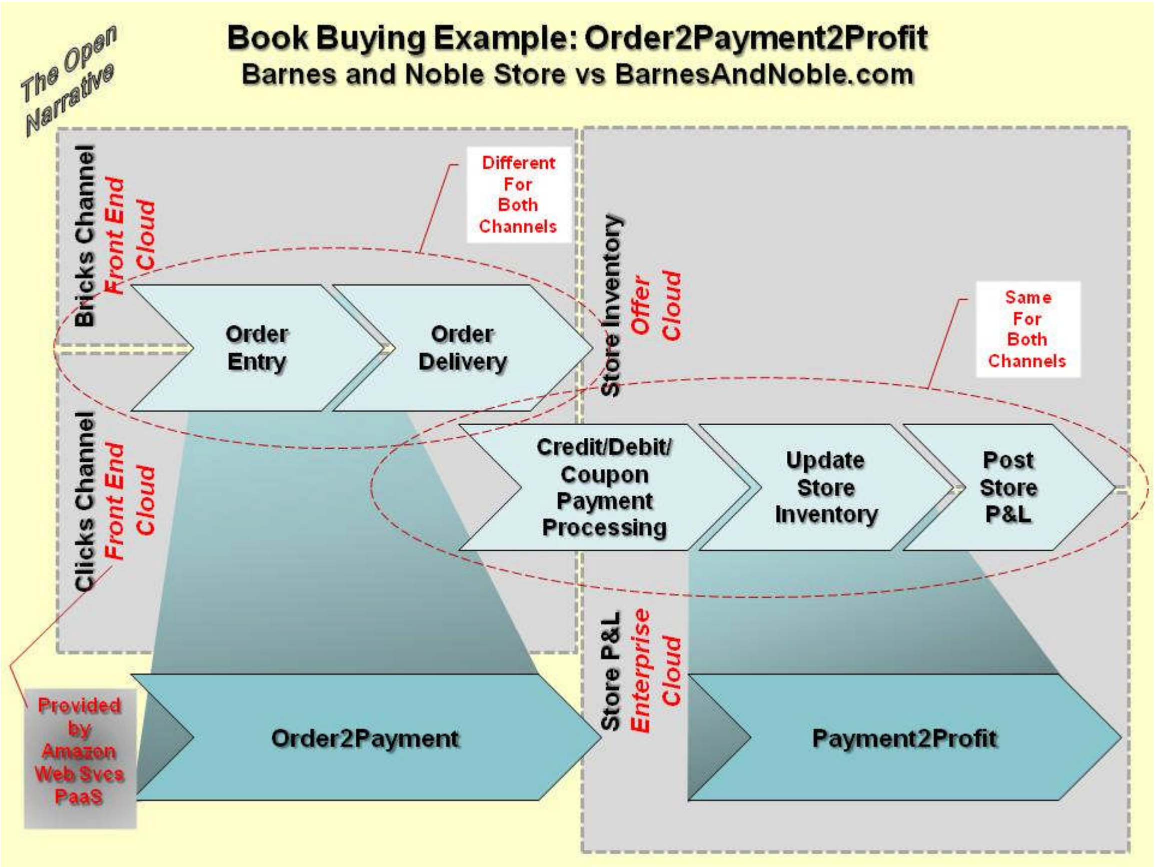

Figure 6: Book Buying Example

Book Business: An Example of Business Aligned Clouds

For example, think of buying a book which can be done in a Barnes and Noble store or through their on-line Web site, the Bricks Channel and Clicks Channel, respectively.

It is interesting to note that Amazon provides a Platform as a Service for Barnes and Noble on-line since 2005. This was the remedy from settlement of an infringement suit by Amazon when Barnes and Noble literally copied the Amazon.com site changing only the branding. This Settlement established the inchoate Amazon Web Services cloud.

Back to the example, putting it all together for book buying, Order2Payment can be realized by a Work Flow serially composed of two flows, OrderEntry and PaymentProcessing.

The example here is one of self-service Order2Payment in either

- **the Bricks Channel experience**

  OrderEntry: enter store; browse shelves; select item(s);

  PaymentProcessing: proceed to checkout queue; pay for item(s) with credit/debit card, cash or coupon; exit store with item(s);

  or

- **the Clicks Channel experience**

  OrderEntry: surf to Barnes and Noble Web site; search for item(s) of interest, selecting into the shopping cart;

  PaymentProcessing: proceed to checkout; pay for shopping cart item(s) with credit/debit card or coupon; exit Web site; await shipment of item(s);

Notice that this is the Barnes and Noble value producing Business Process, Order2Payment with two distinct Work Flows. Although the two component Work Flows, OrderEntry and PaymentProcessing, are named the same and accomplish the same ends, viz., sell items to customers, they are very different with differing implementing Tasks in different Environments (, in this example, Physical Store or Web Merchant Site) with very different security characteristics and requirements.

The Platforms are different, as well, providing different Front-End Clouds, but have a common Enterprise Cloud function, Store P&L, and a common Offer Cloud, Store Inventory.

While deeper consideration of this example would be interesting, this has been an only an illustration of the hierarchy of vocabulary categories. Future work will elaborate these concepts in more depth with respect to the advantages and difficulties of using such precise language.

Cloud Ecosystem Stack

Our operating Cloud vision extends into a 7-18 month Strategic View and Roadmap of how to achieve The Intelligence Cloud. It details an Ecosystem Stack, each layer representing a category to own for participants in the Ecosystem.

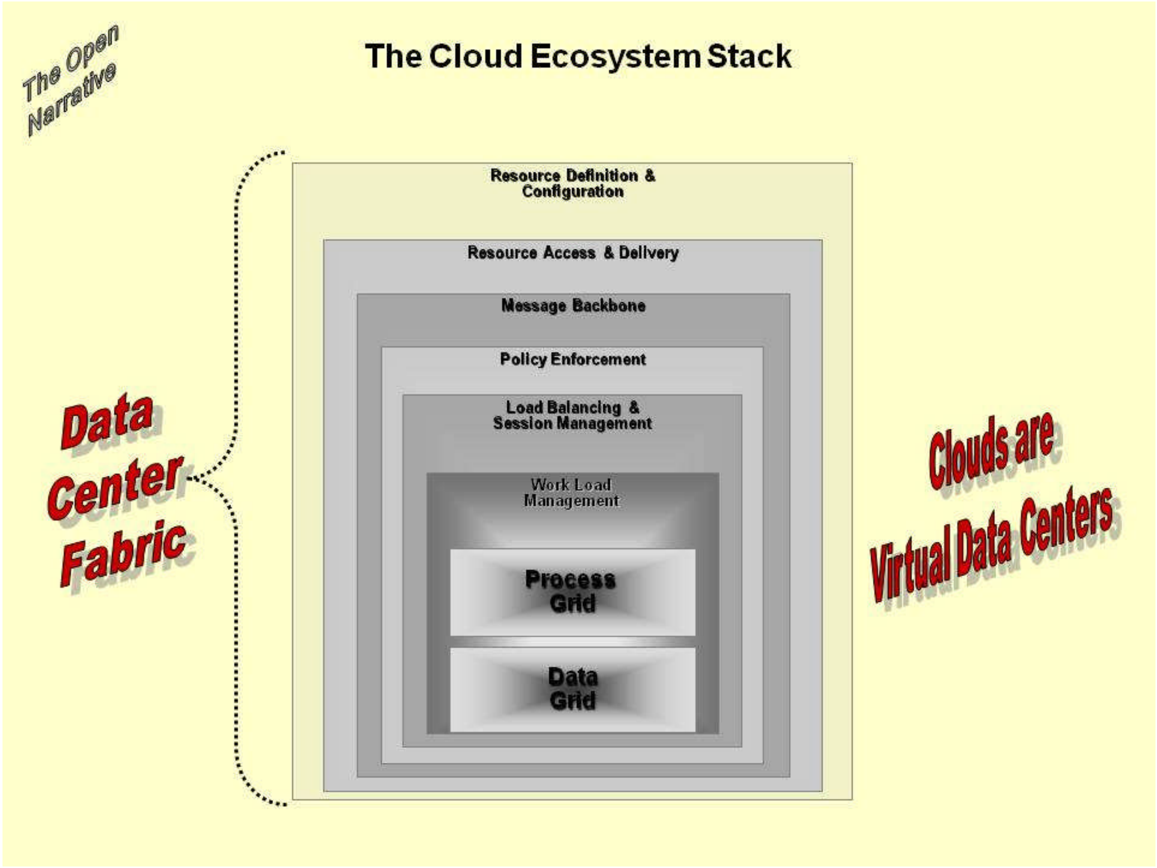
The Open Narrative
The Cloud Ecosystem Stack
Resource Definition & Configuration
Resource Access & Delivery
Message Backbone
Policy Enforcement
Load Balancing & Session Management
Work Load Management
Process Grid
Data Grid
Data Center Fabric
Clouds are Virtual Data Centers

Figure 7: Cross Section of the Cloud Ecosystem Stack

This Ecosystem Stack is seen as defining and refining an Open Narrative on Provisioning Clouds to Business needs. Knowledge and Information Grids put demand on Data Centers while Resource Definition & Configuration down through Process/Data Grid supply that demand at negotiated service levels.

Below each stack layer is discussed in brief.

Process Grid/Data Grid

Process/Data Grid is the core of computing, not just the number of cores on a chip. We still are data centric in our outlook. It is the Bachman ‘68 Turing Lecture: we now pass processes by data, not vice versa. He likened it to the Copernican Earth to Sun centric shift in view of the Solar System. We are in the 40th year of that shift.

Consequently, in Clouds, replicating generic data is more problematic than replicating process images.

The operating concept is a federated virtual cloud for either processors or data.

The Open Source Hadoop (http://hadoop.apache.org/) parallelization works for each in an important use case: static analytics like VaR Monte Carlo Pricing of Risk—about 1-24 hours, an important use case in business intelligence.

Category contenders: Leader—GridGain, Inventigo fastXML, Eucalyptus, Right Scale (Process), Oracle Coherence, Gemstone (Data)

Work Load Management

The High Performance Compute/Data Grid (HPCDG) executes specific functions reliably and consistently. The Work Load (Logical Units of Work—LUns) needs to be scheduled, dispatched and monitored efficiently.

This category is evolving and being removed from the HPCDG layer.

Category contenders: Leader—Platform Computing (schedule/dispatch), Eucalyptus (monitor)

Load Balancing & Session Management

With multiple resource pools able to execute classes of LUns (Logical Units of Work), this category provides upward capabilities to maintain user request knowledge across lower level stateless resource invocations. This is especially true when work is "sprayed across" LUN execution facilities.

Otherwise, state needs to be globally available and consistent which is an intractable if not unsolvable problem in asynchronous, decoupled systems.

Category contenders: Leader—F5 (Load Balancing), Sonoa Systems (Session Management)

Policy Enforcement

Policy enforcement is the centerpiece category driving the Secure Content Aware Network. Security is the prime enabler for Clouds as enterprise grade technology. This category determines the adoption speed at which secure and compliant computing in Clouds can proceed. The Policy Enforcement Category is at Geoffrey Moore's ***Crossing the Chasm*** edge (**http://en.wikipedia.org/wiki/Crossing_the_Chasm**). The issue is coordination with all the other layers of the Ecosystem Stack.

This category needs to lead (pull?) the other Ecosystem category layers across the Chasm with it. Its formulation is the main contribution to the Industry commonweal.

Category contenders: Leader—XML/XACML (IBM DataPower, Cisco Securent. Sonoa Systems), NimbusSecurity Policy State Life Cycle

Messaging Backbone

Clouds involve asynchronous, decoupled systems. Messaging backbones are the effective mechanism to facilitate asynchronous communications amongst decoupled components. Message itself is an abstraction for Service Request, Event Dispatch and Data Stream.

This category includes instrumentation infrastructure of components for monitoring as well.

Category contenders: Leader—XML (IBM DataPower), JMS (http://en.wikipedia.org/wiki/Java_Message_Service)/Open MQ (https://mq.dev.java.net), Open Enterprise Backbone (OEB: http://www.newglobalenterprises.net/docs/EnterpriseServicesBackbone-Std1.pdf)

Resource Access and Delivery

This category includes technologies that (1) virtualize access point connections to Enterprise Resources (Services, Data, Money, Goods) and (2) provision the configurations in which the access points exist.

Included are the downward facing APIs that deliver tools and services on components of interest.

Category contenders: Leader—DataSynapse Fabric Server (service capability), 3Tera (automated provisioning)

Resource Definition and Configuration

This category is the top layer of the Compute Supply Stack. Definition and configuration are driven from Requirements of Business Environments which represent the Compute Demand part of the Stack.

As indicated in the diagram above, this is the outside layer of Data Center Fabric that provides end-user toolkits (upward facing APIs).

Category contenders: Leader— JBOSS SOA Platform (open: www.jboss.com/products/platforms/soa), SOA Software Service Manager (proprietary: www.soa.com/index.php/products/service_manager)

Conclusions & Take-Aways

Current State of the Art

External Cloud is better developed than Internal Cloud because it is somewhat of a green field while Internal Clouds must deal with legacy integration.

External Cloud providers are maturing as viable business operational facilities and have been quite useful in low security circumstances as with Analytics Calculation Facilities.

External Cloud Providers offer walled gardens at the moment, but integration technologies and methods are available to use. Large vendors like the usual suspects are talking Cloud while the Amazons and RackSpace are walking Clouds.

Internal IT organizations are learning about Clouds and what it means to architect internal systems as multi-tenant and multi-landlord.

The Security issues are well under way to resolving and Operational Monitoring is quickly following.

Going Forward with the Long View

We are entering an era when we have both the knowledge and the computing power to really prove things about deployments BEFORE they go into production. This is knowledge and reasoning as the ultimate security infrastructure.

To reason about deployments, more formality is required. But, the constant challenge is to make this formalism accessible. We need “Tools. Tools. Tools.”

A US economic partisan would say, “Who says the US is becoming devoid of tool and die makers?” This virtual manufacturing demands new specialized tools and dies. Now we use “Self-Service” for tool and “Template” for die.

IT is the Tool and Die making of the new manufacturing.