

WEB SERVICES STRATEGY

By

Stephen B. Miles

Submitted to the Alfred P. Sloan School of Management
in Partial Fulfillment of the Requirements for the Degree of
Masters of Science in the Management of Technology

at the

Massachusetts Institute of Technology

June 2003

© Stephen B. Miles

All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly paper
and electronic copies of this thesis document in whole or in part.

Signature of Author

Alfred P. Sloan School of Management

May 9, 2003

Certified by

Charles H. Fine

Chrysler Leaders for Manufacturing Professor of Management Science; Thesis Supervisor

Certified by

Henry Birdseye Weil

Senior Lecturer; MIT Sloan School of Management;

Thesis Reader

Accepted by

David A. Weber

Director; MIT Sloan Management of Technology Program

WEB SERVICES STRATEGY

by

Stephen B. Miles

Submitted to the Alfred P. Sloan School of Management on May 10, 2003 in Partial Fulfillment of the Requirements for the Degree of Masters of Science in the Management of Technology at the Massachusetts Institute of Technology

Abstract

“Everything is connected to everything” [El Aleph \(1945\), by Jorge Luis Borges](#)[1]

This thesis addresses the need to simplify and streamline web service network infrastructure and to identify business models that best leverage Web services technology and industry dynamics to generate positive business results. Web services have evolved from the simple page-display protocol of their origin and now reach beyond the links that simply updated web data dynamically from corporate databases, to where systems can automatically transact. These Web services represent a series of network business technology standards and capabilities that irrevocably change the way in which businesses will do business. In fact, every business today is a networked business and has opportunities to grow using Web services.

This study focuses on the implementation challenges in the financial services market, specifically the On Line Transaction Processing (OLTP) sector where legacy mainframes interface with multiple tiers of distribution through proprietary EDI links. The OLTP industry operates under stringent regulatory requirements for availability and audit-ability of not only who performed what transaction, but who had access to the information about the information. In this environment organizational demands on network infrastructure including hardware, software and personnel are changing radically, while concurrently Information Technology (IT) budgets are under pressure. The strategic choices for deploying web services in this environment may contain lessons for other industries where cost effective large scale processing, high availability, security, manageability and Intellectual Property Rights (IPR) are paramount concerns.

In this paper we use a systems dynamics model to simulate the impact of market changes on the adoption of innovative technologies and their commoditization on the industry value chain, with the aim of identifying business models and network topologies which best support the growth of an “Open Systems” network business. From the results of the simulation we will derive strategic recommendations for networked business models and web services integration strategies to meet Line Of Business (LOB) objectives.

Thesis Advisor:

Professor Charles H. Fine

Chrysler Leaders for Manufacturing Professor of Management Science

Thesis Reader:

Professor Henry Birdseye Weil

Senior Lecturer; MIT Sloan School of Management;

Acknowledgements

I would like to thank my wife Ingrid and son Stephen who put up with my absence during this year of study and supported me in fulfilling my dream to study at MIT. With my warmest thanks and love I dedicate this thesis to Ingrid.

Literally hundreds of people contributed directly and indirectly to the insights in this thesis developed over the course of this year at MIT. I would like to begin by thanking my thesis advisor Charlie Fine, who I first heard speak at a conference on “Innovations in Management: Leadership, Management and Innovation”¹ during my first week at MIT. Charlie’s work in understanding patterns of industry dynamics in “Clockspeed, Winning Industry Control in the Age of Temporary Advantage”[2] and his MIT telecommunications roadmapping initiative have served as a reminder of the challenge of identifying the key drivers in complex systems. The framework of looking at underlying technology drivers, industry structure, regulatory and external influences, with the aim of determining strategy is largely derived from the “Graduate Seminar in Operations Management and Technology Roadmapping” 15.795 cases presented during the fall of 2002.

Thanks also to my thesis reader Henry Weil for his effective use of systems dynamics in competitive strategy analyses, and for sharing his real-world engagements with leading firms. Key questions we ask in this paper are derived from consulting engagement cases in which we participated during Henry’s “Competitive Strategy Course” 15.923 in the spring term of 2003. Henry’s influence on Management of Technology student theses from prior years, most notably David Munsinger’s thesis on the wireless equipment value chain, [3] and Masahisa Kawashima’s thesis on Telecom Value Chain Dynamics and Carriers’ Strategies in Converged Networks [4] was helpful in thinking through the evolution of the web services value chain. I would also like to thank Professor Tom Eisenmann of the Harvard Business School for his teachings in “Managing Network Businesses” Spring 2003 class and “Internet Business Models” text [5] which have furthered my interest in network businesses.

¹ MIT Industrial Liaison Program Conference May 8-9, 2002

The leadership of the MIT Sloan Management of Technology Program (MOT) in the history of technology businesses, starting with founder Jim Utterback, his course 15.353 on “Research Themes” and his work on a Dynamic Model of Process and Product Innovation,[6] has been helpful to me in sorting through the fallout from the collapse of the telecommunications bubble to find new opportunities. Furthermore Jim’s analysis of boundary conditions for academic truisms such as “first mover advantage” and “significant mass” have been useful reminders of the imperative to challenge the assumptions underlying accepted trends of the day, including “web services,” the topic of this analysis! Special thanks as well to Jim’s student and good friend Mark Meyer, Professor of Entrepreneurial Studies at Northeastern University who supported my application to the MIT Management of Technology Program.

On the cover of “The Experience Economy,”[7] Joe Pine, who presented to his former teacher Jim Utterback’s class last fall, there is a picture of the Paolo races in Siena. Informal ties that bring people together such as guilds, neighborhoods, even the colors of their uniforms, lend a rich diversity to the human experience. It is leather workers like those in Umbria who produce the goods which are today sold on Saks Fifth Avenue through Italian fashion houses. A combination of loosely coupled yet tight linkages are required for a business that must, as we learned from David Simchi-Levi in 15.778 Supply Chain Management[8] class last summer, design and manufacture a year ahead of the market. Harold Rheingold describes these network phenomena which have been enabled by new communications capabilities in “Smart Mobs, The Next Social Revolution:”

“Technologies and methodologies of cooperation are embryonic today, and the emergence of democracies, carnival, intelligent new social forms depends on how people appropriate, adapt, transform and reshape the new media once they are out of the labs of the engineers- as people always do.” [9]

In the end, successful societies are ones in which groups of people can form networks dynamically to align themselves on a given purpose, and then reconfigure according to revised priorities.

I owe a special debt of gratitude to the Semantic Switch team which formed in Howard Anderson’s New Enterprises class, and to the groups of engineers who have steadfastly met over the course of two semesters evaluating requirements for XML infrastructure and wireless web services for entries to the MIT \$50K Competition. Special thanks to Raefer

Gabriel, former CTO of True-Exchange, for his hands-on knowledge from building the J2EE back-end infrastructure for the New York Mercantile Exchange, and to Mario Harik, MIT Graduate Student in Engineering, for his willingness to pitch in and code web services using Amazon's API's and/or Photoshop edits as the case required. I would also like to thank the 15.923 Tata Consulting Services Team, and Suresh Sunderrajan and Samir Sanghani in particular, for their work on the Vensim model reflecting the impact of market conditions and sensitivity measurements on key indices. Special thanks also to Tim Panagos, MOT classmate and former CTO of Pegasystems, who stepped in the middle of this project for his work on the CICS web service design which is featured in the case at the conclusion of this study.

Finally, I would like to thank Michael Harte, CTO of PNC subsidiary PFPC, for introducing me to his IT team based on a chance encounter in the hallway at the Sloan Tang Center one crisp early October morning. I later discovered that Michael arrives at PFPC Westborough facilities from his home in Manhattan at the same early hour on every trip. I would further like to thank Per Gyllstrom, PhD, Managing Director of IT, Global Enterprise Architecture for PFPC, for sharing his architectural perspective and for validating the strategic approach to integrating web services proposed in this thesis.

Preface

After fifteen years experience building businesses in the IT and communications equipment sector, this year at MIT has provided me the opportunity to research the technology trends which underlie dramatic changes in global networks. When I enrolled in the Management of Technology (MOT) program my intention was to explore the market for an XML switch to support machine-to-machine communications, including standards being developed at the World Wide Web Consortium for the “Semantic Web.” My prior involvement in “Voice over IP” (VoIP) enabling technologies at NMS Communications led me to ask whether machine to machine XML sessions, similar to a voice over IP calls, would need authentication, QoS, session set-up and tear-down with interfaces to management and billing/charge-back systems. Since IP networks are by definition stateless, technologies for managing synchronous interactions could be applied to keeping these XML sessions “alive.” The product of this research was an MIT \$1K Competition Semantic Switch business plan written with a team of MIT students in Howard Anderson’s 15.390 New Enterprises class in the fall of 2002. How best to manage these machine-to-machine XML sessions to build businesses is the subject of this thesis.

Further evolution of my thinking on network businesses occurred in Tom Eisenmann’s “Managing Networked Businesses” class at Harvard Business School this spring. One of Tom’s guests was the founder of Neopets [10], now the 6th most popular web site. This sparked my interest to try building what David Reed from the MIT Media Lab Digital Life Consortium calls a “Group Forming Network” (GFN). The result was the House of Wisdom MIT \$50K entry to develop a wireless web service based on the GFN model. Advisors included Ted Selkirk, MIT Media Lab, Dr. Kaveh Pahlavan, Director of Wireless Research Center at WPI, and Surj Patel, Imagineer at Orange.

Being in an academic environment this year has given me the chance to step back to analyze the underlying dynamics of a business that I have been immersed in for the past decade. I am deeply grateful for the ideas and input from my professors and colleagues, as well as to the guests who have presented here at MIT Sloan School and at the Harvard Business School. This thesis is my attempt to share some of the learning opportunities that I have been blessed to experience with others.

Table of Contents

Abstract	2
Acknowledgements	4

Preface:	7
Table of Contents	7
1. Introduction and Background	10
BACKDROP	10
OBJECTIVES	11
IS SOMETHING NEW HAPPENING?	12
LEGACY INTEGRATION	14
SCOPE AND SUCCESS METRICS	15
INTERNET TOPOLOGY	17
EVOLVING TRAFFIC CHARACTERISTICS	20
SOCIAL NETWORK PATTERNS.....	21
SCALE FREE NETWORK BUSINESS OPPORTUNITIES AND COMPETITIVE THREATS.....	24
WHAT'S NEXT	29
2. Web Services Definition	31
OVERVIEW	31
WEB SERVICE FUNCTIONS	31
APPLICATION SERVERS	32
WSDL PROTOCOL	33
ENTERPRISE ARCHITECTURE MODEL.....	34
SERVICE-ORIENTED PROCESS	35
TRADEOFF IN SCALABILITY VERSUS HIGH LEVEL ABSTRACTIONS	35
3. Technology Drivers	39
INNOVATION CHARACTERIZATION	39
NETWORK INFRASTRUCTURE TECHNOLOGY DRIVERS.....	40
COMMUNICATION COSTS	41
SERVER CONSOLIDATION	43
DISK CONSOLIDATION A.K.A. "SIMPLER IS BETTER"	45
WEB SOFTWARE SERVICES	46
EDI EXTENSION	46
DATA EXCHANGE.....	48
SERVICE ORIENTED PROCESSES	49
4. Web Services Industry Structure	51
CROSSING THE CHASM	52
TOP OF THE S-CURVE.....	53
DOMINANT PLATFORM	54
APPLICATION SUITES QUADRANT.....	55
BOTTOM OF THE S CURVE:.....	55
BOTTOMS UP" WEB SERVICES MODEL.....	57
CORBA INTERFACE LEGACY	58
ENTERPRISE INTEGRATION APPLICATIONS (EIA)	60
5. Regulations and Standards	65
XML STANDARDS DRIVERS.....	65
xBRL SEA CHANGE.....	70
XML PROTOCOL STACK	70
SECURITY	71
SECURITY ISSUES IN FINANCIAL SERVICES	74
MANAGEMENT	74
INFRASTRUCTURE REQUIREMENT	75
TESTING.....	76

6. XML PROXIES.....	78
HARDWARE XML SERVICES.....	78
SOFTWARE XML SERVICES	81
NEW WEB SERVICES BUSINESSES	82
7. Causal Loop Modeling and Scenarios	84
BUSINESS MODEL.....	84
WEB SERVICES COMMODITIZATION.....	86
DISRUPTIVE APPLICATIONS	87
COMPETITION DYNAMICS FROM INCUMBENTS.....	88
8. Architectural Recommendations.....	89
KEY QUESTIONS	89
DATABASE ADAPTER CONNECTIONS.....	89
PUBLISH AND SUBSCRIBE MESSAGING	89
BUILD & MAINTAIN A PLATFORM INDEPENDENT SERVICE MODEL.....	90
THE CASE FOR SERVICE ORIENTED ARCHITECTURE.....	91
THE CASE FOR 'PRIMITIVES' BASED STANDARD INTERFACES.....	91
ARCHITECTURAL OPPORTUNITIES.....	92
ISV's LEAD THE WAY.....	93
IBM SOLUTION 3 YEARS OUT.....	93
END CUSTOMER WEB SERVICES TRADE-OFFS	94
9. Case Study – ISV CICS Web Service Proposal	95
INTEGRATE ACROSS WHATEVER IS NOT GOOD ENOUGH	95
BUSINESS OBJECTIVES	95
CICS WEB SERVICES ARCHITECTURE	96
END USER IMPLEMENTATION STRATEGY.....	97
10. Business Strategy	99
CORE PROCESS REQUIREMENTS	101
PRESERVING REAL OPTIONS	102
NETWORK EFFECTS	104
A BUSINESS THAT VIEWS ITSELF AS A WEB SERVICE - MVNO	104
COLLABORATION SOFTWARE	106
STRATEGY CHOICES.....	108
FUTURES.....	109
Bibliography Footnotes.....	112
APPENDIX 1 MODEL.....	119
APPENDIX 2. EIA WORKSHEET	122

1. Introduction and Background

BACKDROP

As this thesis is being written, web services are breaking down barriers between disparate systems, organizations and creating webs of new relationships. Value chains are today being “Blown to Bits” as Phil Evans of the Boston Consulting Group described in his 1999 book by that title[11]. Telecommunications, the very industry that spawned IP networks, was the first to suffer this fate in the stock market crash of 2000. Service providers today find themselves relegated to being a dumb transport in a session-less state, having lost track of transactions on the web that are growing in new and different patterns. This thesis makes the case that machine-to-machine transactions, based on maturing web services standards, are powerful enabling technologies that managers can deploy to drive business growth for the coming decade. This thesis offers a strategy for determining the optimal web services architecture and business model for specific business requirements.

Corporate web sites today can be powerful network business tools when web services are strategically deployed to support high yield customized demand chain solutions for particular customers, as well as being tools for employee and supply chain partner productivity. Firms such as eBay and Amazon represent business models that are breaking new ground. Incumbents are also leveraging web services to get closer to their customers on the model of Anoldo Hax’s “Delta Project, Discovering New Sources of Profitability” [12] as exemplified by Dell’s success in integrating around customer configuration requirements all the way back through their supply chain, resulting in growth that generates positive cash flow.² Howard Rheingold documents in his new book “Smart Mobs”[9] how mobile Internet access is allowing even people who don’t know each other to act in concert.

Globalization and the Internet are impacting every corner of the economy. The cost of connectivity has plummeted at the physical layer of the network, while standards for exchanging information at the inter-domain layer have reached a point where machines can be made to communicate. The effects of the converging global network infrastructure are upon us. ERP, Supply Chain and CRM applications extend the reach of the Enterprise up and down the value chain. The *‘lingua franca’* or common language of this exchange is

² MIT Sloan lecture by Michael Dell Chairman and CEO, Dell Computer Corporation September 26, 2002 Tang Center, Wong Auditorium

XML, which provides the alphabet from which words can be constructed. Much like HTML enabled all types of documents to be viewed by web clients anywhere, XML standards for data presentation are evolving to where one database machine can transact based on data from other machines, based on a common ontology. This thesis sets out to trace the underlying network dynamics and to set forth paradigms for network topologies and business models to leverage these effects.

We will examine how the topology of the Internet underpins new business models and opportunities. The underlying pattern of “scale free” networks, which can be visualized as a graph where Y is a linear function of X , as Albert-Laszlo Barabasi of the University of Notre Dame delineates in his new book “Linked, The New Science of Networks,”^[13] promise both extraordinary growth opportunities and vulnerabilities for those acting as the “hubs” of information exchanges. Furthermore, as Nicholas Negroponte, Chairman, MIT Media Laboratories, pointed out in his talk at the 2003 MIT Information Technology conference on “Extreme Communications: A Radical Rethinking of Business Technology and Regulatory Strategies,”³ network technology holds the promise for distributed development as an alternative to mass urbanization of the poor which have flowed into African and South American urban ghettos, for countries like India and China which retain 80% rural populations. Web services technology developments are the subject of this study.

OBJECTIVES

In this first Chapter we identify objectives for this thesis and provide an outline of chapters to follow. This thesis on web services strategy makes the case that new patterns are emerging in the business world, driven by disruptive technology changes that both threaten existing businesses and offer new opportunities to create value. “Scale free” networks promise both extraordinary growth opportunities and vulnerabilities for those who choose to be the “hubs” of these new information exchanges. The analysis of these technology changes and how to build a competitive network business strategy is the subject of this thesis.

³ *Extreme Communications: A Radical Rethinking of Business, Technology and Regulatory Strategies. The 2003 MIT Information Technology Conference.* April 15-16 2003; Wong Auditorium, Tang Center; Charlie Fine and David Clark. Organizers

We will analyze the underlying technology trends supporting web services, their impact on the network value chain, and identify IT infrastructure requirements to support a value added network business model. One way of phrasing the choices that Chief Information Officers (CIO's) must resolve in order to deliver a web services capability is as follows: "Do web services require a new systems architecture, i.e. a new platform, requiring a 'top to bottom' migration plans, or is the technology more accurately represented as a compilation of standards-based interfaces which extend existing systems on the model of legacy Electronic Data Interchange (EDI), only cheaper and easier?" Narrowing the perspective to what an IT manager/consultant might ask: "How do I evaluate the purchase of Enterprise Integration Application (EIA) software and services from one of the major web services applications suite vendors versus developing custom integration to my legacy environment using legacy EDI connections-based protocols?" To phrase the question even more succinctly, from the perspective of a CIO, as Toby Redshaw, Vice President at Motorola asserts in a recent Information Week article, "You better get the architecture right. It might kill you if you don't." [14]

IS SOMETHING NEW HAPPENING?

The notion that new forces are being unleashed in the competitive marketplace that will require businesses to change strategy is not a premise that is universally accepted. This is especially true of skeptics who were on the sidelines of the market boom during the late 90's and who witnessed with some glee the crash of the dot.coms that were to have been models for the "new economy." Economist Stan Liebowitz, in his book on *"Rethinking the Networked Economy: The Real Forces that Drive the Digital Marketplace,"* [15] says that while his studies show peer to peer downloads now exceed the number of music purchases, he does not see a new business model emerging. The author questions whether network effects are any different from traditional economic arguments that "economies of scale" favor monopolies and a "winner takes all" paradigm. While he points to the differences between "weak lock-in" with high switching costs and "strong lock-in" of proprietary format as the source for traditional monopolist strength ⁴ he does not believe that long term 'strong lock-in' effects are possible to create.

⁴ Mr. Liebowitz emphasizes the distinction between two kinds of lock-in. The question of compatibility is central to both. One kind of lock-in arises simply because switching to a new product involves a cost beyond the purchase price: costs of learning how to use it, for instance, or the difficulty of using it alongside products you already own. Mr. Liebowitz calls this self-incompatibility, or weak lock-in. But there is also strong lock-in. This arises if a new product is incompatible with the

The argument that we present sees the emergence of a networked business model, leveraging web services connectivity, with different characteristics from a traditional monopolistic “winner takes all” model. A striking example that serves to illustrate how different models for network businesses can be came to light in a response from Tony Scott, CTO of General Motors (GM), during his visit to the Management of Technology (MOT) program this spring.⁵ Tony was asked which GM business would be worth more in 10 years, the car manufacturing business, which he had described at some length as being a low margin high volume business requiring growth just to cover pension benefits of 2.5 workers to every current employee, or the On-Star business, which is evolving towards a high margin digital two way communications satellite link. Tony’s answer that he was unsure which would come out on top speaks volumes for the promise of the On-Star network business. One might also note that using web services to tighten links to and reduce the number of the 30,000 current GM suppliers could also contribute to a more powerful core business.[17]

Technology implementation issues are significant where customized legacy systems are involved. The web has primarily been used, up to now, as a tool for allowing humans to communicate with machines. When the Internet was designed (i.e. before the Web), it consisted of a number of fairly low-level protocols like sockets, readers, writers, ports, and IP connections. These protocols were designed for machine-machine communication and

choices of other consumers—and if, because of network effects, this external incompatibility reduces the value of the product.

The point is that weak lock-in is very common, indeed pervasive. Many new products have to overcome self-incompatibility. People do not buy a new computer every three months even though the product is improving all the time. Learning to use a new word processor is a bore; for most users, a rival has to be much better, not merely a bit better, to be worth the trouble. Note that if slightly better products are rejected because of self-incompatibility, this is not inefficient: it would be inefficient to buy such a product, incurring all the costs, unless the improvement was big enough to justify it. To repeat, weak lock-in is nothing new.

Strong lock-in is different, because of the network aspect. Strong lock-in means that consumers won’t move to a new and much better product unless a lot of others jump first. If they could somehow agree to move together, they would all be better off. But they cannot. Strong lock-in reflects a failure of co-ordination, it causes economic losses, and in theory it does create opportunities for decisive first-mover advantage. But how common is it, even in the new economy? Mr. Liebowitz is forthright on this. Strong lock-in is not merely uncommon, he says, there is actually no known instance. 16. Economist, T., *The First Will Be Last: Rethinking the Networked Economy: The Real Forces that Drive the Digital Marketplace*. The Economist, Sep 26th 2002.

⁵ 15.379 Seminar in Management of Technology, David Webber Instructor; April 10, 2003

were certainly not easy for humans to use. On top of these low-level protocols were built higher level programs like email and ftp connections, relatively easier for humans but still challenging by today's standards. Then the Web came along, also based on the same low-level protocols, but designed to make it as easy as possible for humans to interact with the Internet. (People who've been frustrated by slow response times when trying to access their favorite web sites may dispute its user friendliness, but compared to where we were a few years ago, it's massively improved.) We now find ourselves in a situation somewhat unusual in the technology field: in a way, the web is now easier for humans to use than for computers.

One would expect that, since all an intelligent agent need be able to do is automatically surf the internet, that it should be able to do so as well as a human. The problem is that the Web is not necessarily an easy place for an automated agent to get around. If it has a question it wants to find, it can probably find a list of related sites pretty quickly through a keyword search (a major accomplishment from where we were a few years ago). But how can it actually get the data it needs out of a web site? The formats that different web sites use are probably as many as there are web sites, and the formats within a given web site often change every few months. What's an intelligent agent to do? The answer lies in creating protocols which make it easier for computers to access data on the Internet, and this is the basic idea of the Semantic Web which we address in Chapter 5.

LEGACY INTEGRATION

We will address the architectural tradeoffs organizations need to make in order to integrate web services with their existing infrastructure. There are still today more transactions processed on IBM CICS & IMS systems, the original modular computing systems, than on the Internet in its entirety. 490 companies of the Fortune 500 generate on CICS more than 30 billion transactions, or \$1 trillion worth of business, each and every day. IT departments in large corporations manage over 30,000 CICS code modules written in COBOL, Assembler or PL/I. The oldest modules are 25 to 30 years old; yet still process mission-critical data everyday. Due to the extraordinary level of customization, these applications endure and will continue to function for many years to come.

Additionally, in recent years, enterprises have supplemented mainframe functionality with similar logic developed or acquired on smaller platforms such as Unix and iSeries (formerly AS/400). Developers built the majority of these applications in the same stovepipe fashion as their mainframe cousins, and therefore these systems display many of the same characteristics. As a result, organizations devote between 70 and 80% of yearly IT budgets to maintaining and evolving this code.⁶ The preponderance of legacy mainframe systems in major corporations is one of the reasons the size of the web services market is expected to grow rapidly according to a recent IDC study:

FRAMINGHAM, Mass., April 2, 2003 - According to a new report published by IDC, approximately 3,300 Web services-based projects were implemented in North America in 2002. IDC's Web services simulation model forecasts North America (United States and Canada) Web services-based professional services spending will break the billion-dollar mark in 2003.

IDC's Web services simulation model forecasts North America (United States and Canada) Web services-based professional services spending will break the billion-dollar mark in 2003 and increase to \$2.7 billion in 2004 - a growth rate of 146% over 2003. The cumulative opportunity between the beginning of 2003 until the end of 2007 will be \$22.2 billion.⁷

Web services technology offers significant opportunities to drive revenue with new one-to-one services and through new channels, to cut costs by integrating further into the supply chain and, more fundamentally, to drive innovation through inter-linkages which have only recently become possible due to the changes in the economics of network technologies that we will explore.

SCOPE AND SUCCESS METRICS

The scope of what one should address in analyzing network business opportunities was clarified for me by Reed Hunt, the former Chairperson of the FCC, when he visited Professor Tom Eisenmann's "Managing Network Businesses" class at Harvard Business School this spring. Reed Hunt invited the class to imagine that we had been appointed to his FCC post. How would we define the FCC jurisdiction in this information age, what would be the scope of communications infrastructure that we would seek to regulate, and what success metrics would we use? The class concluded that overall objectives should include lowering cost of communications in conjunction with some rise in GDP and

⁶ Geoffrey Moore, Davidow Ventures, presentation, Saturday 12/7/02 MIT Venture Capital Conference

⁷ <http://www.webservices.org/index.php/article/articleview/974/1/7/>

insurance of universal coverage. These metrics can equally be applied to the implementation goals of a successful web services strategy.

Reed Hunt declared that using these metrics he created an unmitigated success in the telecommunications industry. As evidence he cited Bill Gates' famous .NET speech stating that Microsoft's software growth would rely on declined cost curve for communications, just as over the last decade Microsoft had relied on the declining cost of computing as predicted by Moore's law.⁸ Just as the initial Microsoft software development model was based on adding complexity and functionality without additional hardware costs, XML is now integral to the next generation collaboration capabilities of Microsoft applications which run over communications networks at a cost that Bill Gates predicts will go to \$0.⁹ Those of us emerging from the telecommunications industry might take note of the declining profit pool and steep declining price/performance curves characterized by this remark.

The scope of networks to be analyzed in this essay therefore includes all transport infrastructure including traditional wire-line and wireless telephone service providers, internet backbone, CATV, broadcast and satellite networks, as well as the Enterprise WANs, and a new class of service providers on the ASP or Amazon model. . We will use as the success metrics of our web services strategy the ability to lower costs in conjunction with some rise in productivity and insurance of universal enterprise participation.

⁸ Bill Gates further describes the challenge today in the following terms:

"In some industries, the issue is not so much faster time to market as it is maintaining time to market in the face of astronomically rising complexity. The Intel Corporation, for instance, has consistently had a 90-day production cycle for its chips, which power most PCs. Intel expects to maintain this 90-day production rate despite the increasing complexity of the microprocessor. The number of transistors in the chip has increased from 29,000 in the 8086 in 1978 to 7.5 million in the Pentium in 1998, and the microprocessor's capability has grown ten thousand-fold over the same twenty years. By 2011, Intel expects to deliver chips that have 1 billion transistors. This exponential improvement stems from Moore's Law, which says that the power of microchips doubles every eighteen to twenty-four months."

http://www.microsoft.com/billgates/speedofthought/looking/chapter_9.asp

⁹ **REDMOND, Wash. -- July 24, 2002** -- Microsoft Corp.'s Chairman and Chief Software Architect Bill Gates today outlined the company's vision and road map for phase two of Microsoft® .NET, the company's two-year-old software initiative for connecting information, people, disparate systems and devices. This next phase continues to build upon the XML-based interoperability of Web services, broadening the benefits to individuals, developers and organizations of all sizes. These efforts encompass software investments to help break down the technological barriers between people, systems and organizations as well as barriers to greater knowledge, trust and everyday use.

<http://www.microsoft.com/presspass/press/2002/jul02/07-24NETDayUmbrellaPR.asp>

INTERNET TOPOLOGY

In order to understand the underlying dynamics driving network infrastructure it is helpful to review the topology of the Internet. We can conceptualize the Internet, or a corporate intranet, as having two distinct levels. At the physical level, the nodes are the routers, and edges are the fiber or cable connections between them. At the inter-domain (or autonomous system) level, we can visualize connections from edge device to a single node representing the entire web and back out to an edge device as in the following schema. [18]




FIG. 1. Network structure of the World Wide Web and the Internet. Upper panel: the nodes of the World Wide Web are web documents, connected with directed hyperlinks (URI's). Lower panel: on the Internet the nodes are the routers and computers, and the edges are the wires and cables that physically connect them. Figure courtesy of István Albert.

Rev. Mod. Phys., Vol. 74, No. 1, January 2002

Figure 1[18]

Network businesses and the web services that run on them rely on Internet and intranet infrastructures that are growing and changing as new applications such as Web services

evolve. A variety of organizations have tracked the growth of the web from both a physical routing path and logical IP address standpoint, as the two examples which follow show:

Skitter data depicting macroscopic snapshot of Internet connectivity: [19]

Map of Internet: colored by IP addresses
By William R. Cheswick: [20]




Figure 2




Figure 3


The first picture represents skitter data depicting a macroscopic snapshot of Internet connectivity, with selected backbone ISPs (Internet Service Provider) colored separately. The second picture shows us which IP addresses are connected. In a seminal article in the IEEE Magazine Topics in Internet Technology section by the title of “Understanding Internet Traffic Streams: Dragonflies and Tortoises,” Nevil Brownlee and KC Claffy come to the following conclusions about the evolving characteristics of internet traffic as they effect network design:

We find that although most streams (about 45% of them) are *dragonflies*, lasting less than 2 seconds, a significant number of streams have lifetimes of hours to days, and can carry a high proportion (50% to 60%) of the total bytes on a given link. We define *tortoises* as streams that last longer than 15 minutes. We point out that streams can be classified not only by lifetime (*dragonflies and tortoises*) but also by size (*mice and elephants*), and note that stream size and lifetime are independent dimensions. We submit that Service Providers (ISPs) need to be aware of the distribution of Internet stream sizes, and the impact of the difference in behavior between short and long streams. In particular any forwarding cache mechanisms in Internet routers must be able to cope with a high volume of short streams. In addition ISPs should realise that Long-Running (LR) streams can contribute a significant

fraction of their packet and byte volumes – something they may not have allowed for when using traditional ‘flat rate user bandwidth consumption’ approaches to provisioning and engineering.

EVOLVING TRAFFIC CHARACTERISTICS

As in the case of using Erlanger formulas to calculate “busy hour calls” for provisioning an SS7 network, this article provides an analysis of network traffic that highlights the “busy hour” volume for Web TCP traffic as well as the evolving packet characteristics:



1(a): Traffic summary at Auckland (UA), kb/s vs time for 5 m intervals

Figure 4 [21]

Since 2000 Internet link speeds increased as users migrated to cable modem and DSL connections, backbone links were upgraded from OC3 (155 Mb/s) toward OC48 (2.4 Gb/s), and ISPs installed newer, faster routers to handle increasing packet loads. At the same time computer hardware improved; systems with 1 GHz processors, 512 MB memory, 20 GB disk drives and ever increasing I/O bus speeds became common. This dramatic increase in network and computer capability has allowed users to work with ever larger files. As a result we now observe that the average size of web objects has increased considerably, with web objects up to 50 kB becoming common. Along with increasing file size, the last few years have seen the rapid growth in usage of an ever increasing set of peer-to-peer file sharing systems. e.g. Napster, Gnutella, E-Donkey, etc.

These peer-to-peer applications have significantly changed the traffic mix, so that a higher overall proportion of their streams have large numbers of bytes. In addition to streaming protocols carrying audio and video programs, VoIP or multimedia conferencing are increasingly common. Clearly these trends will continue.[21]

This difference between the underlying infrastructure and web traffic that plays out on this framework is an issue that must be addressed in planning to match network architecture with new web services applications.

SOCIAL NETWORK PATTERNS

An overlay network to the IP address layer occurs at the social and business level where value is exchanged and transactions occur. Stanley Milgram, a sociologist at Harvard University, in the late 1980's was the first scholar to postulate how tightly connected we are to one another in his book "The Small World." Any person in the world can be traced to any other by a chain of five or six acquaintances. That is, despite the six billion inhabitants of our planet, we live in a "small world." [22] This feature of social networks came to be known as "six degrees of separation" after John Guare's Broadway play and movie. In addition, sociologists have repeatedly argued that nodes (i.e. people) in social networks are grouped in small clusters, representing circles of friends and acquaintances in which each node is connected to all other nodes, with only a few weak links to the world outside their own circle of friends. Further analysis on how close we are to one another has been the subject of numerous publications. Physicist Albert-Laszlo Barabasi in his work "Linked, the New Science of Networks:" describes the history leading up to this discovery.

"Networks and graphs have long been studied in a prolific branch of mathematics known as graph theory. Until recently, the absence of detailed topological information about large complex systems, such as communication networks or cells, meant that networks were modeled as "random graphs". Hungarian mathematicians, Paul Erdős and Alfred Rényi, introduced the most widely investigated random-graph model in 1960. Their influential model consists of N nodes, each of which has a probability, p , of being connected to another node via a link ...

A road map that has cities as nodes and motorways as links is a good example of an exponential network because most cities are located at the intersection of motorways. In contrast, networks that can be described by a power-law distribution look more like the airline route maps found in glossy in-flight magazines. Although most airports are served by a small number of carriers, there are a few hubs, such as Chicago or Frankfurt, from which links emerge to almost all other US or European airports, respectively. Just like the smaller airports, the majority of documents on the Web have only a few links... The physics of the Web Feature: July 2001 Statistical mechanics is offering new insights into the structure and dynamics of the Internet, the World Wide Web and other complex interacting systems. [23]

Both these physical and logical connections turn out to follow a consistent arithmetic progression. Physicist Albert-Laszlo Barabasi and his colleagues at the University of Notre Dame in Indiana in 1998, using a Web crawler, set out to map this connectedness of the Web. They were surprised to find that the structure of the Web didn't conform to the then-accepted model of random connectivity. Instead, their experiment yielded a connectivity map

that they christened "scale-free" which can be seen in the following diagram adjacent to a random model:¹⁰




Figure 5

A subsequent study in "Computerworld", resulting from "a collaboration between IBM, Compaq and AltaVista has subsequently found that the shortest distance between any two nodes in a sample of 200 million is 16. This value is in good agreement with the prediction of 17 for a sample of this size.

These results clearly indicated that Web interactions represent a small world, i.e. the typical number of clicks between two Web pages is about 19, despite the fact that there are now over one billion pages out there. And as Lada Adamic of Stanford

¹⁰ For a proper answer we need a full map of the Web. But, as Lawrence and Giles have shown, even the largest search engines cover only 16% of the Web. This is where the tools of statistical mechanics come in handy - they can be used to infer the properties of the complete network from a finite sample. To achieve this our group at Notre Dame constructed small models of the Web on a computer, making sure that the distribution of links matched the functional form that we had previously measured.

Next we identified the shortest distance between two nodes, defined as the number of clicks required to get from one page to another, and averaged over all pairs of nodes to obtain the average node separation, d . Repeating this process for networks of different sizes using a technique called "finite size scaling" - a standard procedure in statistical mechanics - we inferred that the average node separation is given by $d = 0.35 + 2.06 \log(N)$, where N is the number of nodes. This expression predicts typically that the shortest path between two pages selected at random among the 800 million nodes (i.e. documents) that made up the Web in 1999 is around 19 - assuming that such a path exists. This path, however, is not guaranteed because the Web is a directed network, i.e. a link from one page to another does not imply the existence of an inverse link. Consequently, not all pairs of nodes can be connected - a feature factored into the calculation that leads to the expression for d . 13. Barab*si, A.-L., *Linked : the new science of networks*. 2002, Cambridge, Mass.: Perseus Pub. 280.

University in the US has shown, the Web also displays a high degree of clustering. The probability that two neighbors of a given node are linked together is much greater than the value expected for a random network without clustering. Results from our group indicate that the Internet follows suit - the typical separation between two routers is nine. In other words, a packet of data can reach any router within 10 hops, and the network is highly clustered, demonstrating that the small-world paradigm has rapidly infiltrated the Earth's newly developing electronic skin as well. [24]

What impact will this phenomenon have on the design of web services infrastructure? What we have seen in this section of the analysis is that most exchanges are within a limited number of relevant parties. On the one hand we can deduce from this information that the number of connected parties or systems on a given web services transaction in all likelihood will be > 19 , and probably under 10. On the other hand businesses would like to capture the "scale free" network effect that allows one to concatenate many such groupings into larger network patterns described by Howard Rheingold in "Smart Mobs:"

In the economics of computer-mediated social networks, however, four key mathematical laws of growth have been derived by four astute inquirers: Sarnoff's Law, Moore's Law, Metcalf's Law and Reed's Law. Each law is about how value is affected by technological leverage. Sarnoff's Law emerged from the advent of radio and television networks in the early twentieth century, in which a central source broadcasts from a small number of transmitting stations to a large number of receivers. Broadcast pioneer David Sarnoff pointed out that the value of broadcast networks is proportionate to the number of viewers.¹¹ I asked him [David Reed] what led him to Reed's Law. 'I had the first 'eureka' when I thought about why eBay was so successful'. eBay, which has turned out to be the only hugely successful profitable e-commerce business, doesn't sell any merchandise; it provides a market for customers to buy and sell from each other. eBay won because it facilitated the formation of social groups around specific interests. ..At that time I had been reading Fukuyama about social capital. Fukuyama argues in his book Trust that there is a strong co-relation between the prosperity of

¹¹ On the origins of Metcalf's Law: The often-cited Moore's Law is the reason the electronic miniaturization has driven the hyper-evolution of electronics, computers, and networks. In 1965 Gordon Moore, cofounder of Intel and one of the inventors of the microprocessor, noted that the number of elements that could be packed into the same amount of space on a microchip had doubled every eighteen months in the future. Anything that doubles and redoubles grows large very quickly, from 2,250 elements in Intel's first processor of 1971 to 42 million elements in the Pentium 4 processor thirty years later. Computers and electronic components have driven industrial growth for decades because they are among the rare technologies that grow more powerful and less expensive simultaneously...

What happens when you link devices based on Moore's Law? When ARPA wizards gathered at the Xerox Palo Alto Research Center (PARC) in the early 1970s to create the first personal computers, one of the engineering aces, Bob Metcalf, led the team that invented the Ethernet, a high speed network that interconnected PCs in the same building. Metcalf left PARC, founded 3Com, Inc., cashed out and came up with Metcalf's Law, which describes the growth of value in networks. The math is simple and is based on a fundamental mathematical property of networks: the number of potential connections between nodes grows more quickly with the square of the numbers of nodes.

national economies and social capital, which he defines as the ease with which people in a particular culture can form new associations.

...Human communication adds a dimension to the computer network. I started thinking in terms of group - forming networks (GFNs). I saw that the value of a GFN grows even faster-much, much faster- than the networks where Metcalf's Law holds true. Reed's Law shows that the value of the network grows proportionately not to the square of the users, but exponentially. That means you raise two to the power of the nodes instead of squaring the number of nodes. The value of two nodes is four under Metcalf's law and Reed's Law, but the value of ten nodes is one hundred (ten to the second power) under

Metcalf's Law and 1024 (two to the tenth power) under Reed's law - and the differential rates of growth climb the hockey stick curve from there. This explains how social networks, enabled by email and other social communications, drove the growth of the network beyond communities of engineers and to include every kind of interest group. Reed's Law is the link between computer networks and social networks."[9]

SCALE FREE NETWORK OPPORTUNITIES & COMPETITIVE THREATS

Phil Evans of the Boston Consulting Group, in presenting an introduction to Albert-Laszlo Barabasi's work to Arnaldo Hax and Gerhard Schulmeyer's 15.928 Graduate Seminar on Strategic and Management Consulting,¹² argued that both the physical and logical layers of the internet follow "scale free" patterns, with logical addresses exceeding physical nodes by a factor of five. According to Phil these "scale free" networks of tightly interconnected nodes promise both extraordinary growth opportunities and vulnerabilities for those at the "hubs" of information exchanges.

How important the effects of scale networks are to the design of web services infrastructure can be apprehended in terms of the downside of a network with these characteristics. Scale networks present weaknesses which hackers have been quick to exploit.[25] for example on February 7, 2000, "Mafia Boy," a 15-year-old Canadian boy, overwhelmed first Yahoo and then Amazon, eBay, CNN, eTrade and Excite with phantom web hits. "Love Bug" remains one of the most frequent viruses, a reminder that once overcoming some initial threshold, a virus can be very difficult to track down and root out. [26] Tony Scott, the CTO of General Motors, shared during the MOT dinner on April 10, that his biggest nightmare was of a virus

¹²April 9, 2003 4:00PM Phillip Evans, author of Blown to Bits and Senior VP at The Boston Consulting Group, Adaptable Networks: Aspects of a New Business Paradigm

His three questions are:

Businesses pursue many different possible goals, and networks have many different possible topologies. What kinds of networks are best suited for what kinds of purposes?

As businesses emphasize adaptability, wrestle with deep uncertainty, try to respond to 'weak signals,' what kinds of networks are likely to emerge?

What does this imply for traditional concepts of management, such as the boundaries of the enterprise, external competitive strategy, and internal organization and accountability?

that spread as quickly as the last one, infecting every client in the GM network of hundreds of thousands in a matter of hours. Security and authentication remain daunting challenges as businesses evaluate moving to a net centric model.

For incumbents disruptive technology changes are particularly challenging, as Clayton Christiansen describes in "The Innovator's Dilemma,"[27] in the best tradition of Austrian economist Joseph Schumpeter, who wrote in the first half of the 20th century about his idea that innovation brings "gales of creative destruction." [28]

Schumpeter noticed that technology arrives in clusters -- with electrification come dynamos, generators, transformers, switch gear, power distribution systems; with mass production and the automobile come production lines, modern assembly methods, "scientific management," road systems, oil refineries, traffic control. These clusters, if they are important, define an era. They eventually change the way business is done, even the way society is conducted. [29]

How should an industry like Music and Film respond to Napster's peer-to-peer viral incursion into their intellectual property? Steve Blumenthal, former CTO of Genuity, in my interview with him¹³ said that he witnessed at the Genuity Network Operations Center (NOC) a 50% drop in Internet traffic the day Napster was closed down. By contrast is there new business model that can revive both the communications service provider and the music industries? Low-cost transaction networks together with the ability to create value in network platforms may give birth to new "multiplicative" forms of technological development that we will explore herewith.

While some value chains, such as traditional book distribution and retailing, are getting blown up, other incumbents are using the reach of web services technology to extend their legacy businesses. One need look no further than the supply chains of Wal-Mart and General Electric, some of the largest companies in the world, which reach far into China. The business model Amazon used to emerge as the largest bookseller represents more than a forgiving capital market in boom times. Two hundred computer engineers at Amazon, many of them from MIT, have built a prototype for a new business platform which is based on web services and personalization engine technology which supports "selling everyone everything" according to my interview last fall¹⁴ with Amazon CTO AI Vermulen, co-author of "The Elements of Java Style," [30]. In addition to Amazon's direct reach and providing outsourced

¹³ Meeting at Au BonPain, Kendall Sq. April 4, 8:30AM.

¹⁴ Conference call with AI Vermeulen, CTO Oct 20 12:30PMEST;referred to Rob Frederick, Director of Web services, Amazon

web hosting for companies including Toys R Us and Circuit City, Amazon offers an open web service accessible via XML API's ¹⁵which over 900,000 independent retailers have linked to.

An example of one of the most successful social network service from the telecommunications industry was the MCI "Friends and Family" offering, a clear example of Fukuyama's social capital at work,[31] which cost \$50M and two years to develop and became one of the most successful service launches in history.[32] Nextel's "Push to Talk" service, whereby groups of contractors can communicate concurrently, has had the highest Average Revenue Per Subscriber (ARPU) in the industry. It is no wonder that Motorola just purchased Winphoria who purport to have a "push to talk" softswitch service capability. [33] Arnaldo Hax illustrates a series of similar examples in his chapter "Delta.com: Reinterpreting the Internet History." [12]

Examples of scale networks, and network effects in general, are the subject of some of the most popular web sites¹⁶ and businesses school classes today. Tom Eisenmann's HBS class #1760 "Managing Network Businesses" this spring concentrated on this topic. Covering a large range of information industries we identified network dynamics including straight network effects (Metcalf's law – addition of another fax machine...), complementary effects (Gameboy Game Cartridges) and liquidity effects (more participants make markets less volatile). We studied common traits of successful network platforms that include reinforcing loops and high switching costs. Vendor ratings, as exemplified by eBay, had the effect of producing reinforcing loops, as did editor ratings from Epinions, and "karma points" from Slashdot publishing. The CEO of NeoPets presented the business case in class for a site which now ranks #6 in the world for stickiness. An example of a successful network business can be seen in the following picture:

¹⁵ API <http://associates.amazon.com/exec/panama/associates/ntg/browse/-/1067662/103-5864171-7785428>

¹⁶ ECONOMICS OF NETWORKS This server was created by Nicholas Economides neconomi@stern.nyu.edu.

At this site, you will find a collection of information on economic issues of networks, such as the telephone and fax communications networks, the Internet, financial exchange and credit card networks, as well as on "virtual networks," such as the virtual network of all Windows or all Mac computers. *Ranked as one of the top 5 economics sites worldwide by The Economist magazine.*




Figure 6¹⁷

I have since learned that creating a network effect that “sticks” is harder than it looks. MIT AI Labs postgraduate Luis Bettencourt who created this illustration, in trying to model the probabilities for success, an exercise which he equates to predicting fads or fashions, came to the following conclusion: “prediction of definite [network effect] outcomes where there is no substance is very very difficult.”¹⁸ As Andrew McAfee and Francois Xavier Oliveau point out in their article in the MIT Sloan Management Review [34] network effects can be slowed or even reversed by such phenomena as saturation, noisiness (“cacophony”), and Economist Paul Krugman’s insight that “The flaw in Metcalf’s Law ... becomes apparent: Most people

¹⁷ A social network with clustering (into 4 communities) and a hub.

¹⁸ -----Original Message-----

From: Luis Bettencourt [<mailto:lbmet@ins.mit.edu>]

Sent: Thursday, February 13, 2003 1:29 AM

To: Stephen Miles

Subject: RE: Watts's colloquium

> Obrigado! - de nada ... perhaps you can tell me more about your project. Otherwise it is difficult to know where to start.

People like narrative, a certain amount of simplicity and animation. And everything else being equal will do what their neighbors are doing - most people hate being left out of the loop. But it all depends - kids really like repetition and simplicity, like Sesame Street and the blue dog show, what ever the name is now (blues clues something or rather).

Ah - Had a look at neopets - I bet what happened was that these pets with their intricacies (stories, possible narratives) became popular with some kids in school, then they all will try to find what the next pet is going to be and do to stay ahead of their peers in the info loop. And find the pet that can beat the hell out of their friend's pet, etc. The site also offers "Never fear: latest gossips, news ..", - bingo - jeez it's tough to be a kid these days - there's so much to know. This place is basically the arms dealer for the popularity arms races of 9-13 year olds. In any case, as i stressed, prediction of definite outcomes where there is no substance is very very difficult.

Best,

Luis

Luis M. A. Bettencourt

<http://www.mit.edu/~lbmet>

Center for Theoretical Physics, 6-316]

have nothing to say to each other!”¹⁹ It is nonetheless clear that over this same network topology viruses spread at a furious rate that we easily recognize from their biological brethren.

In his presentation referred to earlier, Phil Evans of the Boston Consulting Group spoke about the paradigm change that is occurring in businesses. Scale Free networks, uncovered in the analysis of Internet connections where 99% of routers have less than 10 ports,²⁰ have parallels in business that he proceeded to outline. Using the example of the pharmaceuticals industry, Phil asserted that small biotech companies, closely linked in Intellectual Property rights (IPR) sharing agreements in “scale free” networks with each other, consistently outperform “walled garden” Research & Development groups at “Big Pharma.” Phil linked the biotech innovations with the creativity of 30,000 Open Source Linux programmers, organized into core and assistant developers, who have generated 30M lines of code, (this represents an investment equivalent to \$3.4 Billion) and are outperforming Microsoft’s core operating system group in speed and reliability.

Phil cited a variety of network effects, ranging from the Philippines population who overthrew of Marcos communicating via Short Message Service (SMS), to the bloggers who brought down Trent Lott, as examples of successful networked modes of organization. Similarly “loosely connected” communities may range from groups of parents concerned about autism (for which Phil presented a picture of linked clusters of web sites) and extend to on-line gamers. Apparently 40% of Koreans play Lineage, while the average Everquest user amongst 1M is already on-line 20 hours/week. Web portal exchanges draw values for Avatars that may reach \$18 (this price multiplied by the number of artifacts created would give Everquest a GDP larger than the country of Iceland).

Phil went on to defend the ideas from his book “Blown to Bits” in the context that at that time major corporations were ignoring the impact that an Amazon might have on the ‘bricks and mortar’ incumbent such as Barnes and Noble. Partially in response to this threat, major corporations did indeed catch on. As a result a wide range of web services are being implemented in large corporations, the trend which gives rise to this thesis in how best to streamline web services development and simplify the associated infrastructure.

¹⁹ <http://www.redherring.com/mag/issue55/economics.html>

²⁰ lumeta.com/graphics/wired.gf; Govindan, Ro and H. Tongmnarkit

Competitive analysis of wealth creation opportunities from these “free scale” networks goes beyond the traditional Porter “Five Forces” strategic planning process that is enterprise centric,[35] to a world that is network centric, where processes can become services. Combining the “richness” of a value added service with the “reach[11]” enabled by web services can create a powerful engine for new business. These are key insights we would like to take from the analysis of underlying network topologies to build a web services strategy that concerns the balance of this thesis.

WHAT'S NEXT

Chapter 2 defines basic web services functionality and identifies different network topology strategies.

Chapter 3 describes key technology drivers and uses frameworks from the Management of Technology program at Sloan to characterize web services innovation.

Chapter 4 presents two distinct analyses of web services. An argument is made for the maturation of web services that is evidenced by 70% market share by top 4 vendors. A subsequent analysis shows that a different approach to web services through point extensions to existing architectures.

Chapter 5 describes the impact of regulations and standards evolution on the growth of web services.

Chapter 6 describes alternative hardware and software XML proxy solutions to support machine to machine communications

Chapter 7 uses causal loop modeling based on the systems dynamic model created by Forrester and as described in John Sterman’s “Business Dynamics, Systems Thinking and Modeling for a Complex World, ”[36] to model different web services strategies in differing Old Business and New Business market conditions.

Chapter 8 takes the results of the sensitivity analysis to make architectural recommendations for which approach to take, depending on business objectives.

Chapter 9 is a case study where we propose a radical simplification to legacy integration infrastructure.

Chapter 10 concludes with a framework to evaluate organizational and Line of Business (LOB) applications requirements.

My hope is that this thesis will encourage organizations to evaluate different web services strategies to grow their businesses. In conjunction with this thesis, technical due diligence is underway on a radically simplified web services architecture that works as a CICS transport service. Such a topology would represent a breakthrough in IT systems design for legacy environments seeking to capitalize on web services while minimizing layers of middleware.

2. Web Services Definition

OVERVIEW

Web Services streamline communication by providing a language-independent, object-model independent set of messaging protocols, programming standards, and network registration and discovery facilities to expose business functions to authorized parties accessing the Internet from any Web-connected device. These web services that run over public and private networks are breaking through geographic, organizational and market barriers, introducing new competitors at every turn. While some have questioned the extent of the impact of the Internet and web services on business, this study supports findings that the pace of globalization and computer-to-computer communication is increasing and that the way in which businesses communicate is changing. We will study how underlying technology changes impact business exchanges.

“By the end of the decade, more than 60 percent of the computer communications will be computer-to-computer. Computer-to-computer vastly speeds up the pace of business. For example, end-to-end supply chains can be automatically adjusted by point-of-sale computers directly communicating with warehouse computers, which in turn directly communicate with manufacturer computers, and, again in the chain, manufacturers’ computers directly communicate with their supplier computers. In addition, computer-to-computer communications can track demand and adjust logistic systems to automatically direct product to geographical points of demand.”
[37]

We begin this section with an introduction to web services functionality. Much like HTML enabled all types of documents to be viewed by web clients anywhere, XML standards for data presentation are evolving to where one database machine can interpret and transact based on data from another machine.

WEB SERVICE FUNCTIONS

Web services can be organized into the following functional categories:

1. Web Services as basic query/response using XML structures to present formatted data
2. Web Services for data interchange, i.e. next generation EDI using WSDL to transact
3. Using Web Services as middleware to create network-aware applications with the equivalent of Remote Procedure Calls (RPC)

The definition from Oracle's web site summarizes core web services constructs:

*"Web Services consist of a set of messaging protocols, programming standards, and network registration and discovery facilities that expose business functions to authorized parties over the Internet from any Web-connected device. More simply, a Web Service is a discrete business process that (i) Exposes and describes itself – A Web Service defines its functionality and attributes so that other applications can understand it, and makes this functionality easily available to other applications; (ii) Allows other services to locate it on the Web – A Web Service can be registered in a sort of electronic "yellow pages," so that applications can easily locate it; (iii) Can be invoked – Once a Web Service has been located and examined, the remote application can invoke its service using an Internet standard protocol; and (iv) Returns a response – When the service has been provided, the results are passed back to the requesting application again over the same Internet standard protocol. There are two types of Web Services - simple and complex. Simple Web Services provide basic "request/response" functionality, are typically not transactional in nature, and provide simple HTTP-S/SSL based security. They are developed supporting three primary Internet standards - SOAP, WSDL, and UDDI. Complex Web Services can be characterized as multi-party, long-running business "conversations" that involve sophisticated security, such as non-repudiation and digital signatures, as well as business-to-business collaboration and business process management. They are developed supporting ... primary Internet standards - ebXML and RosettaNet."*²¹

Next we examine the application server middleware which supports query response, the software standards for the exchange of information and creation of network aware applications.

APPLICATION SERVERS

Gartner defines applications servers as performing the function of web services platform middleware, with most of the elements needed to deliver "Greenfield" web services applications:

²¹ Oracle 9i Web Services Strategy

Java 2 Enterprise Edition (J2EE) application servers (such as WebSphere, WebLogic and iPlanet) are the primary candidates for the role of a Web services provider platform, especially for the new "green field" implementations of Web services. Application servers are the modern implementation of platform middleware ... which also includes such categories of products as TP monitors (such as CICS and Tuxedo) and transactional ORBs (such as Orbix and COM+). Historically, application servers have often served as the basis for enterprise implementations of SOA.

The application server acts as a container to application software. Thus, application servers are the natural containers for the WSDL proxies and the new implementations of Web services. With support for XML parsing and SOAP channel listening, an application server can fulfill all basic requirements of a Web services provider platform. All leading application servers in their current versions are equipped to act in this role.

Application servers impose a programming model on applications. J2EE, COM/COM+, .NET, CORBA and Application Transfer Manager Interface (ATMI) are examples of such programming models. Application servers support only that application software which follows its designated programming model. Thus, the interface proxy programs, to be maintained in an application server, must be written or generated to the particular application server's programming model. While most modern programming models are standard (such as J2EE and CORBA), each application server displays some subtle differences in its implementation of the standard. For practical purposes, a Web services proxy program must be generated specifically for the application server that will act as its container. This assures complete matching of versions, protocols, and various flags and tags that, together, represent the Web services standards. Connectivity between a client proxy and server proxy generated for the same Web service (but by different generators) will often be problematic. Thus, for all practical purposes, the ability to generate WSDL proxies (client and server) is another essential requirement for a Web services provider platform. [38]

WSDL PROTOCOL

Application servers that are best equipped to act as Web services provider platforms will have the tools to generate WSDL proxies as follows:

- Request/reply and one-way messaging access
- HTTP, SMTP, FTP and other Internet protocols
- Non-Internet protocols (MOMs and ORBs)
- In the future, for alternative wire protocol standards, based possibly on ebXML

Messaging or other then-prevailing extensions to SOAP

For practical purposes, a Web services proxy program must be generated specifically for the application server that will act as its container. This assures complete matching of versions, protocols, and various flags and tags that, together, represent the Web services standards. The business purpose of the proxy program is to leverage standard interfaces to core business processes to “get close with the customer.”[12]

ENTERPRISE ARCHITECTURE MODEL

The key to web services is the ability to design a custom interface to both the inside and outside world. A model for how to think about the architectural issues surrounding web services which frames the underlying business processes has been proposed by Gartner.

New formulations of IT architecture are attempting to supply this needed new language. Gartner's Enterprise Architecture Model (see Figure below) depicts IT architecture in terms of four related, increasingly complex layers. The top two layers are for business consumption, while the bottom two are strictly IT-specific and therefore should remain internal to IS. The top layer depicts the emerging virtual business model, and is appropriately termed the "Power Grid." The second layer addresses how the business goes to market, i.e., its business processes. [39]




Figure 7

This architectural notion of web services acting as the presentation layer for a series of processes and services is reinforced in a recent ZapFlash research note on “Service-Oriented Process: Thinking about Processes before Services.”

SERVICE-ORIENTED PROCESS

Business process methodologies are not particularly new. However, what is unique about Web Services-based SOAs, is that in an SOA, a process is a service. Individual services can be composed into a process flow that itself can be described as a service using the same Web Services standards. The ability for services to be consumed by processes that can then be described as services is a very potent notion that shouldn't be overlooked. The services that are consumed by a process might be services local to an enterprise or available at a remote location. In this regard, a Service-Oriented Process exhibits the location independence as well as implementation independence features of loose coupling. A business process description (or flow model) can be published as a Web Service, described via WSDL and interacted with using SOAP. As a result, it becomes impossible to distinguish whether an invoked service is an atomic service or a service composed with business process! . Users simply invoke the Web Service, which has been either specified statically, or determined dynamically through the use of a service registry.[40]

Several trade-offs must be evaluated as we think about how best to architect a series of business processes that we want to make available to the enterprise as a web service. What the Gartner model misses is the tradeoff that is implicit between a high level of generalization and the transaction readiness of working at the level of system “primitives” in a given system. The tradeoffs between scalability versus high level user interface designs are explored in the following article:

TRADEOFF IN SCALABILITY VERSUS HIGH LEVEL ABSTRACTIONS

Building scaleable solutions for complex query and exchange systems has a number of challenges.

Effective management of content (security, transaction management, versioning, query processing, *etc.*) becomes dramatically more difficult as the model (relational, object-oriented, compound, hierarchical, *etc.*) becomes more complex. Resource allocation and time constraints are **not** the only reason the management functions/applications available with relational databases are not available with the more complex content management systems that manage XML, documents, web sites, and other far more complex irregular models. At some point, architects have to choose which facet of their content management system is going to shine and which will struggle to meet minimum requirements.




Figure 8

The illustration above charts the evolution of the content management as it has expanded its information modeling functionality and information management functionality. It is a useful comparison because it shows that the market has overwhelmingly preferred trusted management functionality to enhanced modeling capabilities. The then dominant network (CODASYL) DBMS vendors ridiculed the early relational products (you couldn't even model an organization chart let alone a parts assembly with the relational model!). The buying public decided that elegant modeling could reduce coding efforts and improve design, but management facilities were essential to running a business. Many years later, the early Object-Oriented (OO) DBMS vendors had visions of conquering the then multibillion-dollar DBMS market due their revolutionary information model. Again, the market spoke unambiguously in favor of predictable and scalable platforms.

The support for XML inside repositories continues this tradition. XML is clearly a data type that must (and is) supported. However, a few vendors have tried to use DTDs (or more recently XML Schema) as the data definition language (DDL).²² [41]

In an initiative led by Benjamin Grosph at the MIT eBusiness Center, “we give a new fundamental knowledge representation formalism: a generalized version of Courteous Logic Programs (CLP), which expressively extends declarative ordinary logic programs (OLP) to include prioritized conflict handling, thus enabling modularity in specifying and revising rule-

²² Of course, consumers don't like to choose, they want it all, and the RDBMS vendors have moved to incrementally enhance its modeling capabilities. First came enhanced data types, including data blades and data cartridges from Informix and Oracle respectively. Next came numerous interfaces, filters and data processors that provided persistent storage and various levels of validation of OO-based and XML content. Most recently, the tight bundling of application servers and file system and hierarchical storage management support has been bundled into the DBMS. These developments have seriously cut into the functionality of the basic CMS system. As such, one can see CMS vendors responding by accelerating application enhancements to minimize the importance of these now commoditized (low value) capabilities.

sets. Our approach to implementing CLP is a courteous compiler that transforms any CLP into a semantically equivalent OLP with moderate, tractable computational overhead.”[42] We will come back in chapter 7 to this trade-off between high level object representations versus scaleable transaction performance as it applies to constructing an implementation strategy for web services in a specific business case.

As compared to the Gartner Enterprise Architecture model I have found a more useful framework from an Information Week article last fall which is from Carlson Travel.[43] This framework clearly illustrates the separation between business process and the infrastructure layers to be implemented.




Figure 9.²³

The term service-oriented architecture describes a method for building a corporate software infrastructure that allows different applications to exchange data and processes regardless of the operating systems or programming languages underlying those applications. In this model, an application or a portion of an application is a service that another application or person can subscribe to without extensive custom coding.²⁴

²³ <http://www.informationweek.com/story/showArticle.jhtml?articleID=6504188&pgno=2>

²⁴ Research firm ZapThink that specializes in web services markets released on Thursday its prediction that Service Oriented Architectures will become the dominant designs for networked business systems by 2006. Analysts at Gartner are also predicting that service-oriented architectures

One of the clear benefits of a SOA is that such an architecture helps companies get more value out of existing resources, by wrapping legacy applications in Web Services interfaces and then making those Services available and discoverable on the network. A second thrift benefit that SOAs provide is that they facilitate heterogeneous IT environments. Instead of "rip and replace," moving to an SOA means building bridges between different systems and applications, rather than throwing them out. [45]

This architecture supports the alternative of leveraging Open Source software, as was the case with the Oracle approach illustrated in the first part of this chapter, where an Oracle database was connected to an Apache web server running the LINUX operating system. Whether programmers who have taken that step wouldn't also move to an open source database such as MySQL is a question that is not a question that we will address directly in this essay.

However if web services is in fact the technology to drive the next generation of business, i.e. if Bill Gates bet on .NET and XML is right, and more business applications will be web centric, then the question that any IT organization might ask themselves is how best to support these new business initiatives. "Is web services a revolutionary new platform architecture, requiring a new software architecture and a redesign of Enterprise Systems, or can web services be simply implemented as an evolutionary set of interfaces upgrading legacy EDI and CORBA applications or component object modules?" How do we integrate web services into our object oriented methodologies and the XML embedded in Adobe and Microsoft client applications such as Word into an enterprise architecture? What impact might these changes have on our business model?

will enter mainstream usage this year as per the following report. 44. *software architectures on rise*, in *CNETNews.com*. February 20, 2003.

LaMonica, M., *Flexible*

3. Technology Drivers

INNOVATION CHARACTERIZATION

Stepping back from specific web services topologies for a moment, a helpful model with which to look at the impact Web Services innovation on an organization can be found in Jim Utterback's book on the dynamics of innovation. [Figure 1 A framework defining innovation from "Mastering the Dynamics of Innovation" James Utterback, Harvard Business School Press, 1994]. The model addresses important issues in managing technological innovations regarding the natural locus of innovation, the most appropriate type of innovation, and the array of barriers to innovation. How we master the dynamics of web services innovation is a direct outcome of where we position web services with respect to the core business model of a given business.




Figure 10²⁵

From this diagram one interpretation is of web services as an Architectural Innovation, whereby "The essence of an architectural innovation is the reconfiguration of an established system to link together existing components in a new way." The opportunity with web services under this scenario is to harness and optimize internal systems resources. We can also use this diagram to position web services as a case of Modular Innovation capable of overturning existing value chain relationships external to the company, which is made possible through increasingly well defined XML standard interfaces. Examples of

²⁵ UTTERBACK, J., and ABERNATHY, W. [1975], "A dynamic model of process and product innovation", *Omega*, Vol 3, No.6, 1975 Pergamon Press, GB pp. 639-56.

modular innovations come to mind from the MOT tour of the Computer Museum in Silicon Valley this spring with Gordon Bell, former DEC CTO, who expounded on the design of the first modular computer, the IBM System/360. “Modular systems are more difficult to design than comparable connected systems. The designers must know a great deal about the inner workings of the overall product or process in order to develop the visible design rules necessary to make the modules function as a whole.” Extending this analogy to the web, the fabric of fiber and switch/routing/wireless infrastructure now circumvents the globe, all the while legacy systems remain connected via antiquated point to point links. How to adopt a web services strategy to match these infrastructure dynamics is the topic of this chapter.

Putting network technology in a broad context, an historical analogy for where communications networks are today may be found in the case of railroads at the turn of the last century when a standard gauge was agreed to, when more railway lines were laid than would ever be used, when steam engines were built in sufficient quantities, so that efficient management of the infrastructure was paramount. “The railroad companies devised management practices to deal with their own complexity and high fixed costs that deeply influence the second wave of industrialization at the turn of the century” as Carlis Baldwin and Kim Clark point out in “The Power of Modularity, Managing in an Age of Modularity.”[46] We will explore how web services can be employed to leverage business opportunities both as an Architectural Innovation and/or as a disruptive Modular Innovation. How does one evaluate whether architectural or modular interface approach is the optimal model, and under what circumstances is this the case?

NETWORK INFRASTRUCTURE TECHNOLOGY DRIVERS

The emergence of Web services as a fundamentally different way of connecting the pieces has a lot to do with increases in productivity, but there are other drivers. Real-time computing, collaborative business, and security are business imperatives at many companies, and all require changes to underlying IT infrastructure. New technologies such as IT appliances, server blades, and grid computing have an impact, too, as does the spread of open-source software and alternatives to conventional storage media. We have organized the key price/performance drivers underlying network infrastructure changes as follows:

- 1) Changes in infrastructure costs
 - Network communication costs
 - Server consolidation

- Network infrastructure requirements
- 2) Fundamental standards for information display
- Centralized
 - Decentralized
 - Issue of modeling functionality versus effective transaction management

COMMUNICATION COSTS

The basic underlying physical dynamics of the network infrastructure tells us something about how technology will disrupt existing network architectures. It is a well recognized phenomenon that communications costs have dropped with respect to capacity increases even more quickly than processor costs in recent years.




Figure 11 DIAGRAM OF CPU AND DATA RATE SPEED INCREASES

How to compare price/performance of network transport to centralized applications, as in the case of Storage Networks - i.e. what is the significant relationship between processing power units and transmission speed - has yet to be rigorously defined, with the result that a valid business model for web-based storage is still in flux.²⁶ Nonetheless when chart the cost of communications against the trend in disk drive performance which exceeds processor growth then we get a series of price/performance curves which resemble the diagram on the

²⁶ Per conclusions from "Live Vault Project Report," Spring Term 15.399 E-Lab; Ken Morse and Barbara Bund, Instructors

next page.²⁷ In this diagram, Red would represent the rate of increase in MBps communications cost, Green represents disk storage MB/\$ cost and Moore's law MIPS/\$ would be the constant in this case, i.e. communications costs are projected to drop at a greater rate than Moore's Law:

Performance Trends





Figure 12²⁸

The impact of the drop in network services costs was brought home to me during the MOT visit to Oracle in March where John Dolan, Vice President of Wireless & Voice Strategy, said



27

[<http://www.nap.edu/html/esi/fig1-1.jpg> DRAM has expanded faster than cpu MIPS/inch²

²⁸ ei.cs.vt.edu/~cs4984/external2/sin_creeping.htm

it is now more cost effective to store email and voice mail files centrally and distribute them, as against processing messaging in a distributed network of Microsoft Exchange Servers. The Microsoft .NET strategy is another example we have cited of the shift to a network centric architecture. This insight was further validated in a recent email that I received which was circulating in a major European wireless carrier:

“Bill Gates said in this speech in New Orleans a couple of weeks ago.. “Mobile is our major focus now.” You can almost feel the presence of his rifle scope on the back the heads of certain companies...

Basically his vision is that all phones will be fully data enabled with transparent connectivity between 802.11 and your mobile carrier... and no more carrier voice mail ... but microsoft's vmail... and no writing to carrier specific api's -- Microsoft will build all voice service including location based api's directly into the OS. etc... ie there's no reason to have to write to Vodaphone or Orange only api's... well only if you don't want your application to ever run on Symbian or java handsets....”

The challenges associated with designing infrastructure to support applications in this new environment are non trivial, as exemplified by the competing requirements for streaming media from sources at the edge such as Windows while wanting to assign QoS priority to latency-sensitive applications such as VoIP and web services sessions which must be kept “alive” to complete a transaction. We will investigate infrastructure platforms to help us manage these sessions in the Strategic Recommendations section of this report.

SERVER CONSOLIDATION

Building on the dropping cost of network communications, companies are increasingly consolidating their server operations to minimize Operations, Administration and Maintenance (OA&M), the cost for which can exceed the price of the original equipment. Open Source Linux platforms running on server blades and clusters of disks are increasingly common, promoted by companies including IBM and Oracle to counter Microsoft's increasing market strength in mid tier markets. A well-known advocate for Linux Servers, Ron Katz, CEO of Orbitz, presented at the MIT Sloan 50th Anniversary Celebration in the fall of '02. In a follow-up telephone interview with his CIO Kevin Malover, Kevin explained how Orbitz broke past previous business records, achieving \$1.5B in sales in the 1st full year of operations. Orbitz' travel scheduling algorithms from ITA of Cambridge, Mass run in a pooled cache memory architecture that stores data downloads from airplane reservations mainframes the world over for access over the web.

Google's systems architecture has further commoditized hardware infrastructure required to support web services and searches, while powering one of the most successful network business models today:


"Mr. Schmidt's announcement was even more far-reaching than that. In essence, what he said was that Google, the hottest young company in technology, had committed the ultimate apostasy: it had declared its independence from Moore's law. Not completely, of course, because that would be impossible, but a selective independence. Often forgotten is that Moore's law has three variables: **price, density, and performance**, each of which contributes to the 100 percent improvement the law promises every couple years. But in practice, technology leaders have always stuck with the density track, adding ever more computing power to their products for the same price. Thus, each new generation of Pentiums and Athlons inevitably leads to new generations of PCs, routers, and game machines. Nobody gets excited that 8088 prices have now fallen to, say, 50 cents, making cardboard box cameras 12 cents cheaper--after all, that's 1974 technology. But Mr. Schmidt's announcement changes all that. "We aren't interested in getting maximum power for a high price," he says. "What we're looking for is maximum functionality and that's a whole different thing." Each of Google's thousands of motherboards (a computer's main circuit board) are designed for the quick switching of components. Even the power supply is held on with Velcro straps: if it burns out, it can be replaced quickly. Recently, when the expensive top-end disk drives used by the motherboards proved inadequate, Google tossed out thousands and replaced them with cheaper, better models. "It's all about what works," says one Google technician."
<http://www.redherring.com/insider/2003/02/moore021003.html>

Nonetheless mainframes in the large financial institutions remain the undisputed champions of transaction processing as compared to consolidated server farms, as Chuck Gallant, Senior Vice President of E-Commerce for PFPC reminded me in our meeting earlier this year.

²⁹ Figures from Gartner Research bear out this assessment, and may even be optimistic according to colleague Raefer Gabriel³⁰

²⁹ January 2, 2003; Meeting at 101FederalStreet, Boston re: XML switch functionality

³⁰ From Raefer Gabriel, CTO: "I am not sure how confident I am that people will be really using J2EE application servers to handle 5000 TPS OLTP needs within 18 months. In fact, true 500 TPS is pretty tough - see, for example, the Middleware Company's benchmarks here (which are generally interesting and useful to check out - they do a big comparison of .NET and J2EE here). Their example based on the J2EE Pet Store example (but highly optimized) gets about 50 TPS on a quad processor server with one J2EE app server, and 18 TPS with another - realistic numbers from my experience, after lots of optimization. See the document here: <http://www.middlewarecompany.com/documents/j2eedotnetbenchmark.pdf> However, many more poorly architected J2EE apps get much, much worse performance than this on even heavy-duty hardware. See this thread for example: http://www.theserverside.com/home/thread.jsp?thread_id=7730 The original J2EE Pet Store example code (this comes with J2EE, so lots of people use it for testing) get around 10 TPS on quad processor PIII/Xeon class machines. Slow as a dog - there's really no way to scale that kind of stuff up to 5000 TPS."



Source: Gartner Research

Figure 13

Gartner Bottom Line: J2EE application servers are rapidly maturing and improving in scalability and performances. However, developing large, systematic OLTP applications is still a technically complex [requiring a] highly systematic effort ... Clients should evaluate the risks and benefits of adopting stand-alone J2EE application servers for large OLTP projects. Should the risk be too great, in comparison to the anticipated advantages, users should instead aim at a combination of application servers (to provide the “global class” features) and traditional but viable platforms, such as CICS, IMS or Tuxedo (to sustain the transactional load). Mainstream users [47]

Where Semantic Switch \$1K MIT Competition had initially focused on building a J2EE platform to migrate mainframe applications to, there clearly needed to be a shift in focus if we were to tackle transaction volumes on this scale.

DISK CONSOLIDATION A.K.A. “SIMPLER IS BETTER”

Disk consolidation has occurred in conjunction with clustered CPU architectures. Randy Rettberg, member of the MIT AI Laboratory and former CTO of SUN Storage Systems, spoke at the MIT Laboratory for Computer Science on the subject of “Life at the Top of the “S” Curve.”³¹ He advised that state-of-the art disk clusters are using simpler disk drives in the latest clustered disk configurations from Sun Microsystems. Randy pointed out that optimal

³¹ April 2nd MIT LCS talk on “Life at the Top of the “S” Curve,”

designs avoid virtualization and attendant multiple levels of firmware releases. One can avoid incompatibilities of firmware releases in a cluster by selecting disk drives with no firmware; ie, cheap and simple ESD drives vs. highly performant, expensive and complex, Fiberchannel and SCSI disk drives in the cluster. The essence of Randy's observation becomes the underlying architectural principle of this thesis... 'simpler is better'.

If it is in fact less expensive to centralize data in one place, more performant to store centrally and transmit remotely than to compress and virtualize data, then one would expect enterprise architectures to change dramatically in the coming years in line with current pronouncements by Microsoft chairman Bill Gates. In the wireless initiative that Bill Gates launched in New Orleans, he pitched replacing voice mail in the Central Office of the phone network with Microsoft Exchange servers.[48 1:30 PM PT #56 1:30 PM PT #56] In the subsequent example from Oracle, Apache running on Linux is Oracle's answer to providing the functionality of a Websphere or .NET web services application suite.

WEB SOFTWARE SERVICES

Web services deliver business functionality that can be categorized as follows, in the context of the three-point framework introduced earlier in this paper:

1. Web Services as next-generation EDI
 - a. Integrate legacy systems to J2EE application servers for display on the Web (Figure 1 below)
2. Web Services options for data interchange
 - a. Provide external API to access data from other business applications
 - b. Interface to new platforms internal or external to the organization (Figure 3 and Nokia example below)
 - c. Expand functionality; wireless extensions would be one such extension
3. Using Web Services as middleware to create network-aware applications with Remote Procedure Calls (RPC)
 - a. Integrate directly into third party applications and sites on supply chain or customer facing sides of the business.

We can visualize these processes in the following diagrams from Oracle:

EDI EXTENSION




Figure 1: Standards for Simple Web Services

Figure 14

Identifying elements from the picture, the functionality of each layer can be described as follows:

Web Services Description Language (WSDL) provides a standard format to describe services interfaces.

Simple Object Access Protocol (SOAP) defines a standard request/reply and fire-and-forget, XML-based protocol to invoke a remote service.

Web Services Description Language (WSDL) proposals address the issue of developing composite services.

Universal Description, Discovery and Integration (UDDI) defines a standard set of interfaces to a repository storing Web services metadata.

Proposed WS Transaction, WS Coordination and Business Transaction Protocol (BTP) standards attempt to define transactional models for distributed services.

An example of a web services interface to new platforms is Nokia's wireless web service solution for the 6300 phone, similar to counterparts from Qualcomm BREW and the Microsoft Mobile Initiative as was presented by Mr. Niklas Savander, Senior Vice President, Nokia Mobile Software at his recent MIT talk.³²

³² Media Tech, Sloan presentation Tuesday, April 1st, 2003




Figure 6. The key opportunity of Web services interfaces in the mobile context lies in open standards server-to-server interfaces enabling more efficient service and content deployment.

Figure 15³³

This area is especially active with Microsoft tackling Nokia partner Symbian's web services JAVA based handset architecture.

DATA EXCHANGE

Initial installations of web servers involved queries from web clients to legacy systems, with a subsequent wave of more complex web services applications involves databases talking to database, smart agents, and an infrastructure of building blocks including Web Portals, Web Servers, EIA capabilities, and Application Servers, linked to the database(s) for transactions:

³³ http://www.nokia.com/downloads/solutions/mobile_software/Web_Services_white_paper.pdf




Figure 3: Standards for Complex Web Services

Figure 16

SERVICE ORIENTED PROCESSES

Instead of thinking of Web Services as a collection of interfaces to software functionality that must somehow be made to connect to other such interfaces, enterprises should approach Web Services as enabling a fundamentally process-driven architecture that leverages distributed processes in addition to distributed services. Distributed services is like distributed computing, only it is based on the notion that service interfaces should be standardized and service descriptors and registries should be used to allow for runtime binding.

The traditional mindset that needs changing is the view that Web Services are an extension of the component object model. To many developers, Web Services are simply "another interface to a compiled object." As a result, they apply their traditional component object design methods, deployment technologies, scalability and reliability approaches, and even terminology. The result: point-to-point implementations of Web Services that are every bit as brittle, tightly coupled, synchronous, and fine-grained as their object-oriented predecessors. In essence, these developers are taking a "bottom-up" view of Web Services where the service implementation is the center of the world. As these developers are quickly realizing, thinking of Web Services as connected to a particular object model (.NET or J2EE, for example) or implementation approach (application servers vs. messaging middleware, for example) makes meeting the basic principles of SOAs a challenge. [40]

In a Service Oriented Architecture scenario, integration goes from being a troublesome chore that must be accomplished through implementing increasing layers of complicated and expensive technologies to a side effect of process execution. The mere act of orchestrating and choreographing a Service-oriented business process accomplishes the task of most

integration goals.

4. Web Services Industry Structure

Charlie Fine in his book "Clockspeed"[2] uses the diagram of a double helix to suggest that horizontally organized industries may revert to vertical integration as customers and vendors consolidate, which in turn is disrupted by new market forces. This industry model would be a likely candidate for example to describe trends in the telecommunications industry as revisions in FCC regulatory statutes enable incumbent LEC's to bundle out their Competitive Local Exchange Carriers (CLEC's).³⁴ We have already covered the impact of regulation on common carrier communications businesses in the introduction to this thesis. At the 2003 MIT Information Technology Conference "Extreme Communications: A Radical Rethinking of Business, Technology, and Regulatory Strategies" which Professors Charlie Fine and Dave Clark organized April 15-16, one of the organizing topics was around the regulatory structure as it effects licensed and un licensed spectrum, with 802.11 emerging as a disruptive winner. As new disruptive technologies kick in, the technology road mapping process changes in step. How we should view the web services industry structure is the subject of this chapter.

The recursive process of industry consolidation and fragmentation was highlighted for me by Allen Carney, Vice President of Marketing with my former employer NMS Communications³⁵ when he asserted that the telecommunications industry, due to regulatory changes, never got past the tipping point to a horizontal structure with standards based interfaces. A competitive marketplace for modules with standardized telecommunications interfaces and improved performance never fully materialized for NMS. In this case the double helix metaphor for this industry might be more accurately depicted as a simple pendulum that swings back and forth, periodically rotating a full 360 degrees in one direction or falling back the other direction.

By contrast the financial services industry structure has a well-established history of transacting over a plethora of legacy EDI networks. Examples of successful business models in financial services where different elements of the value chain are broken into horizontal segments include services that are based on standard interfaces between legacy systems, such as is the case with the Fidelity outsourced contract with Bankers Trust for \$11B stock index fund management, and where State Street operates pension custody services for the Fortune 500 enterprises.

³⁴ <http://www.tiaonline.org/policy/filings/96-149-reply.cfm>

³⁵ 3/17/2003 Framingham MA; Allen is former Vice President of Marketing, Lotus Corporation

CROSSING THE CHASM

We begin by analyzing web services industry structure by identifying technological innovations and where they sit on the technology adoption curve. The question still lingers in many mainframe operations as to whether web services technology is a new and unproven industry, or an extension of existing applications, a modular innovation, or some combination thereof. In this chapter an argument is made for a “top down” view of web services of a consolidated web server market where 75% market share has been captured by the top 4 vendors. A subsequent analysis however shows that, depending on the business goals and objectives, a “bottoms up” view may be preferable, through extensions of existing architectures and/or Open Source solutions. From the Management of Technology core curriculum we have learned about the importance of locating web services technology on the “S” curve of market development as our choice of strategy will differ according to positioning on the bleeding edge, rapid market uptake or decline.

To use Geoffrey Moore’s term from the MIT Venture Capital Conference in February, a recurring question that occurs in interviews on the subject of stability and maturity of web services can be phrased as: “Has web services ‘crossed the chasm’ of adoption?” Are web services over the early adopter challenges and ready to grow at a tornado rate, or has the market for Web Services Application Suites matured (70% are purchased from top vendors: IBM, BEA, Oracle, Sun, and now Microsoft) and started to level off?




Figure 17³⁶

³⁶ : Enterprise Forum Satellite Broadcast June 2, 2002; Greg Erman MarketSoft Corporation

Where would you place web services with respect to the “S” curve of a technology adoption cycle? From the study of the history of technology start-ups in Jim Utterback’s 15.353 course on “Thesis Themes” last fall, we recognize that clusters of new companies occur from the earliest section up to top of the adoption curve until consolidation occurs. Are web services still in a phase where the number of start-ups is proliferating, or is rapid consolidation well underway? In an unexpected result, our analysis finds both trends occurring simultaneously, i.e. that web services are both a maturing technology and one in which new applications continue to emerge.

The general choice a customer faces as to whether to purchase a bundled web server application suite and set of services from a brand name software vendor, or whether to build code to industry standard interfaces using a variety of ISV tools or Open Source software (such as Apache Tomcat) is a dilemma we will explore in this chapter. These are strategic choices which dictate how a business would structure web services implementations.

TOP OF THE S-CURVE

As web services technology matures, we would expect innovation to shift from individual specialized practitioners and customized point solutions to the centralized product development groups with product and process systems tightly interwoven. To the extent that web services represent a package of well-established interfaces and functionality, implementation of these applications will require applications expertise. As Professor Utterback and Abernathy illustrate in their work, “the locus of innovation shifts with the stage of development from earlier disconnected efforts to a more systematic stage:

During the unconnected stage in the development of a process, innovative insight comes from those individuals or organizations that are intimately familiar with the recipient process, rather than those intimately familiar with new technologies... Later, in the systematic stage... these needs lend themselves to context technological solutions and the innovator will frequently be the one that brings new technological insights to the problem. Finally, in the systematic stage resistance stems from the disruptive nature of innovation relative to current practice, i.e. the model begins to help clarify the changing nature of barriers to change.

The view of Web Services as a mature technology platform spearheaded by major systems vendors is an example of where a technology burst onto the scene as a premium high end product, beginning with BEA WebLogic and followed by Sun, IBM WebSphere and now Microsoft’s .NET. The introduction of a standardized interface for inter-systems

communications was led from the top, with high-end software packages implemented by the Big 3 consulting firms and IBM. These new products would have more in common with the introductions of supercomputers, high end workstations (Silicon Graphics...) and other higher order disruptions than with the cheap technologies Christiansen highlights in "Innovator's Dilemma." [27]

"The appearance of a dominant design shifts the competitive emphasis in favor of those firms – large or small – that are able to achieve greater skills in process innovation and integrations with more highly developed internal technical and engineering skills." [6]³⁷ Once critical mass is reached, the number of vendors is expected to drop dramatically now that majors have 70% market share [IDC footnote], partially through a series of acquisitions to include Enterprise Integration Application (EIA) functionality.

DOMINANT PLATFORM

In conformance with the findings of Michael Cusumano in his book on Microsoft, a dominant set of players led by Microsoft, Intel and Cisco have set the stage for the computer industry of late. Here we see the emergence of an equally compelling dominant design for Web Services. IBM, BEA, Sun and Oracle and now Microsoft own 70% market share, with .NET surging into third position as noted below.

IDC has predicted Web services will become the dominant distributed computing architecture in the next 10 years and the value of the Web services opportunity in the US will be USD21 billion by 2007, peaking at USD27 billion in 2010.


IDC's latest report on the topic also pointed out that although only 5 percent of US companies had completed Web services projects by the end of 2002, around 80 percent will have some type of initiative in this area underway by 2008.
<http://www.theregister.co.uk/content/archive/29257.html>

From this perspective Web Services have in fact reached a critical point of consolidation where customers can safely choose Web Services Applications Suites between a limited set of well-established players as indicated in this Gartner Quadrant:

³⁷ "Mastering the Dynamics of Innovation," James Utterback, Dominant Designs and the Survival of Firms page 30

APPLICATION SUITES QUADRANT

That there are general purpose Application Suites providers who have established market leadership is clearly illustrated in the following Gartner diagram:



Source: Gartner Research

Figure 18 Web Services Applications Suite contenders [49]

The application platform suite is viewed in this context as an mature market. Through 2005, customers will face a choice between one-stop shopping with suites and assembling solutions with best-of-breed applications.

BOTTOM OF THE S CURVE:

By contrast we will also want to look at the technology disruption being driven by new web services applications. Using XML, a multiplicity of applications with specific functionality for integration into legacy environments have developed, as have extensions to legacy applications by the applications vendors themselves. We have seen several years of

innovations in the creation of industry specific XML interfaces which build on legacy EDI networks to allow the transfer of data from one system to another, along with the creation of increasingly complex middleware based on object oriented technology.

Suddenly new competitive alliances are being reached that were heretofore impossible. What do you do when your largest competitor bundles your firm out of listing a data feed from a major provider? This was exactly the scenario presented to the team by a large financial services company which instigated this study.

The way we might visualize this state of affairs resembles the series of concatenated “S” curves replacing one another as shown in Rebecca Henderson’s work on technology strategy in the semiconductor manufacturing equipment industry,[50] or Clayton Christianson’s analysis of the sealed media disk drive industry,[27] where each new entrant overcomes the incumbent on the prior S curve, before being overtaken itself by the next generation of disruptive technologies:



Figure 19³⁸ Successive technological disruptions.

One of the points that jump out of this picture is that companies rarely make the leap from one technology to the next. If web services are viewed from the standpoint of evolving functional areas like Applications Servers, Enterprise Integration Application (EIA) servers, etc., then where we are on the broader technology trend impacts what purchase decisions we make of web services software. The critical benefit of a standardized interface is that we should be able to mix products and services from different vendors. This is the promise that the World Wide Web Consortium is holding out against proprietary “de facto” standards from major vendors.

³⁸ http://www.abinavinnovation.com/uploaded_imgs/img23.gif

The disruptive competition for web services suites is based on standardization of data interfaces and types of specific applications. Examples of these functional subsets of web services, some of which overlap, but each of which can represent one of the little “S” curves above, are displayed in the following table which link to vendors in each category via XML:

<input type="checkbox"/> App Server for Web Services	<input type="checkbox"/> Best Book
<input type="checkbox"/> Framework for Web Services	<input type="checkbox"/> GUI for Web Services Product
<input type="checkbox"/> Integrated Services Environment	<input type="checkbox"/> Portal Platform for Web Services
<input type="checkbox"/> Service-Oriented Architecture	<input type="checkbox"/> Web Service Security Solution
<input type="checkbox"/> Web Services Automation Tool	<input type="checkbox"/> Web Services BPM Engine
<input type="checkbox"/> Web Services IDE	<input type="checkbox"/> Web Services Integration Tool
<input type="checkbox"/> Web Services Legacy Adapter	<input type="checkbox"/> Web Services Management Tool/Platform
<input type="checkbox"/> Web Services or XML Site	<input type="checkbox"/> Web Services Platform
<input type="checkbox"/> Web Services Testing Tool	<input type="checkbox"/> Web Services Training
<input type="checkbox"/> Web Services Utility	<input type="checkbox"/> XML Database
<input type="checkbox"/> XML Parser	<input type="checkbox"/> XSLT Processor
<input type="checkbox"/> Most Innovative Application of XML	

Figure 20 ³⁹

“BOTTOMS UP” WEB SERVICES MODEL

The key business opportunity that a “bottoms up” approach to web services enjoys is the ability for a service to be generated as an extension of an existing architecture rather than as one more overlay, requiring new security, authentication and management. The whole category of TP Monitors and Transactional ODB’s present opportunities for development of open, web services components. This finding is supported by a Gartner study of non-standard middleware:

³⁹ <http://www.sys-con.com/webservices/readerschoice2003/>

Mainframe and other TP monitors are the previous-generation, nonstandard platform middleware. Like the modern application servers, they can and will be used to host and expose Web services interfaces. Much of the activity in exposing Web services from TP monitors will be driven by modernizing and Web service-enabling enterprise legacy applications. The leading TP monitors (CICS, IMS, Tuxedo) already support or, by 2004, will support the Web services provider platform-required features (0.8 probability), including a SOAP listener, an XML parser, a WSDL proxy generator for the respective TP monitor platform and the proxy access to the TP monitor's legacy application programming interfaces.[38]

Some of the functional categories which are represented in the above table, the small “S” curves we are visualizing, have been rolled up into larger categories through acquisition, as was recently the case in the buyout of Praja last fall by Tibco Business Factor business activity monitoring software.[14] In each case the customer is faced with the choice of adding a higher layer of middleware or reverting to more of a point solution.


CORBA INTERFACE LEGACY

The beauty of standards based interfaces is that customers can integrate products from different applications. CORBA interface standards were a precursor to XML and were often used to transfer information between different databases. Programmatic integration servers provide a mechanism to package defined sequences of character-based data into programmatic representations, such as Component Object Model (COM), Java or Extensible Markup Language (XML), to supply legacy information to new composite applications. As with their presentation-oriented siblings, programmatic integration servers provide the ability to:

- Design logical flows of transactions
- Invoke those transactions through their character-oriented user interface
- Properly transform and map data, moving from step to step in the flow

For large-mainframe-based legacy environments, programmatic integration servers can provide the steppingstone for application modernization and SOA transition during the next three years. I concur with Gartner's assessment that: “The programmatic integration server market represents a great opportunity for growth. As enterprise demands for integration mature, more-architected solutions are required. Products in this market provide greater flexibility to enterprises looking to expose legacy systems through more-sophisticated composite applications.”[51]

Figure 1
Programmatic Integration Server Magic Quadrant for 2003



Source: Gartner Research

Figure 21

All the while some companies are purchasing general purpose application suites, another set of vendors are pitching scaled-down versions of web services platforms. This emerging sector is one which Gartner has named the Enterprise Service Bus (ESB) category that is emerging as a major market force in 2003, including such products as Iona's Inferno.

Besides Iona, a number of smaller companies are developing relatively stripped-down integration software designed around Web services standards, including Simple Object Access Protocol (SOAP), and Web Services Description Language (WSDL). These ESBs are simpler, low-end alternatives to more comprehensive, but expensive, integration middleware, the report said.

Iona is testing its Inferno product with some of its customers and plans to sell the initial version in the first quarter of 2003, primarily targeting businesses looking to solve "medium-weight" integration problems, Newcomer said. The company will add to the product throughout the year with advanced features such as quality-of-service

guarantees that ensure specified response times and document translation services that convert incompatible XML-formatted data.

The more robust follow-on features, such as quality of service and business process workflow, will draw on Iona's experience with middleware that adheres to the CORBA (Common Object Request Broker Architecture) standard, Newcomer said. CORBA is a set of programming specifications that define communications between applications and components. ...

With these more restricted applications the question a user faces is how to architect these solutions into the larger enterprise software model for the company. Here again the beauty of a "bottoms up" services oriented approach to web services creates a new service which can be published to the organization without incurring additional overheads.

The commoditization effect of many low-cost Web services applications cuts into the market for more comprehensive middleware products from companies like See Beyond, Tibco, and Vitria, as many projects will not demand the full range of features.⁴⁰ Essentially we can see the impact of the double helix cycle as new more specialized applications disrupt their more comprehensive brethren resulting in the classic Christensen disruptive effect, and as platform vendors integrate functionality from mid tier players as in the following example:

"Web Logic's strategic move to include WebMethods-like E-enterprise Integration Application (EIA) is a recognition that much web services work is and for some time will continue to take place in the spaghetti-like world of enterprise application integration (EAI), not the stratosphere of online services. Support for integration means simplified EAI that will help BEA win business from IBM, whose own integration is driven by Global Services."⁴¹

In this scenario we see how specialized new entrants are disrupting their applications precursors, while this functionality is being added by the major platform vendors, whether through acquisition or in-house development.

ENTERPRISE INTEGRATION APPLICATIONS (EIA)

We now come to the issues that confront most businesses when they seek to make their core processes that are locked in legacy environments available as a web service. The single-largest installed base of legacy applications resides on mainframes. Web services can be used to expose business logic developed on these systems. However, even though Web services are meant to isolate the implementation from the access, this technology still is

⁴⁰ http://news.com.com/2100-1001-978608.html?tag=fd_top


⁴¹ <http://www.theregister.co.uk/content/53/29859.html>

more easily understood in terms of the newer application paradigms —.NET or Java. To utilize Web services, IS organizations must spend time and money understanding their established systems and use tactical, programmatic integration solutions.

The overall category of integration software is referred to as Enterprise Integration Applications (EIA) and are generally composed of the following elements:

- Transport and connectivity, including the use of messaging, Internet access, file transfer, database gateways, and Web services standards such as HTTP, Simple Object Access Protocol (SOAP) and Universal Description, Discovery and Integration (UDDI) registry
- Transformation (including syntactic conversion and semantic transformation for XML and other data types)
- Intelligent routing (including content-based routing and publish-and-subscribe)
- Business process management and human workflow
- Adapters and collaborations that tie into other software platforms, packaged applications and industry-standard message sets • Message warehouse and business activity monitoring features
- Ease and power of the development tools
- Throughput, scalability and reliability
- Administration and management

While software products are centered on an integration broker, EIA suites include adapters, management tools, development tools, message-oriented middleware and business process management in some cases. Gartner organizes the niche versus the EIA application suite vendors in the following market quadrant:



Source: Gartner Research

Figure 22.. Application Integration Vendors – Gartner

Gartner is very clear to say that the choice of applications software should be driven by the business need and vertical market requirement rather than a 'one package suits all' approach.

The best solution for a particular job may come from a Niche Player vendor that happens to offer just the right combination of features. Although most of these products can be used for a variety of applications in a variety of industries, most vendors are stronger in a few particular industries or application types than in others. Indeed, some vendors concentrate on a single vertical market to maximize their marketing effectiveness and their ability to deliver a stronger product solution for that industry. It is impossible for even the largest vendors to have a balanced and deep set of broker features, application-specific adapters and domain knowledge across all potential applications.[52]

Web services platforms tend to focus solely on providing Web services capabilities, which result in a more robust and complete implementation. Because they can all be deployed into existing J2EE application servers, the benefits of an application-server implementation of Web services--clustering, load balancing, database connection pooling--also can be

achieved. The disadvantage with this approach is one of administration and support. A single-vendor setup often is more desirable in a complex environment, but the flexibility offered by platforms may be more performant.

Novell got itself a winner with its acquisition of SilverStream's Extend. Rebranded as Novell Extend, this application server offers both the corporate-class functionality required of an enterprise application and the features we'd expect for successful deployment of Web services. ...

Web services CAN BE an extension of existing systems that allows integration into third party applications and portals, providing back end data access and compatibility. Furthermore object extensions which are open as in the case of Systinet and Firestar Software offer the promise of moving data from the IBM 390 and CICS transport and batch processing functions to real-time straight-through transaction engines based on J2EE, .Net... [53]

In summary where we are on the S curve of web services adoption is different depending on the application and industry. The emergence of Websphere and .NET as the dominant platform for web services is unmistakable. Nonetheless, IT management will recognize the fallacy of adopting general-purpose solutions for highly specialized and performant applications. Furthermore the divide in the larger enterprise software market between open systems and packaged software vendors is replicated in web services, with Apache Tomcat leading the way in cutting edge applications and low cost scaleable solutions.

From a core technology standpoint, a helpful chart identifying the major XML evolution from Zapthink follows. The issues of Security, Management and Transactions that are critical impediments to adoption are covered in the next chapter. Principle opportunities for new technologies lie in the area of Dynamic Binding to static Services according to this chart, which we pursue in the context of a proposed legacy web services integration strategy.

ZapThink Web Services Roadmap




Figure 23 ⁴²

Stepping back from the specifics we can see the evolution of aspects of web services technology as both part of a larger “S” curve of well established interfaces and application suites and as a collection of many smaller point solution “S” curves concurrently. The web services application suites we recognize today have largely absorbed the functionality of web portals and cache memory pools, applications servers, transaction and workflow messaging infrastructures and Enterprise Integration Applications. By contrast specialized applications continue to evolve for specific point solutions, depending on the application and legacy infrastructure.

⁴² <http://www.zapthink.com/flashes/07012002Flash.html#topstory>

5. Regulations and Standards

XML STANDARDS DRIVERS

By standardizing interface definitions XML provides the basis for integrating different systems. Customer demand is driving the adoption of these standards, sometimes starting with government regulations dictating standard interfaces for large purchases.⁴³

As Stuart Madnick points out in his work in a recent MIT working paper, while it requires a high level of expertise to implement deployable J2EE solutions, "...one of the most important benefits is that XML does help to create structured web pages, compared with HTML." [54]

Feature	HTML	XML
Extensibility	Fixed set of tags	Extensible set of tags
Tag purpose	Tags describe presentation	Tags describe data content
Views	Single presentation	Multiple views of same document (by XSL)
Orientation	Documents	Documents plus semi-structured data
Search	Keyword search only	Keyword plus field-sensitive queries


Figure 1. Comparison of HTML and XML

Figure 24

Before looking in more detail at industry specific regulatory and standards implementations it may be useful to put the evolution of XML standards in relationship to the content management industry where XML has emerged from as a form of extension itself. As illustrated in this web services market map by Zapthink, we can see how web services coupled with content management rely on a security and management infrastructure as well as the development of Value Chain applications and Industry Vertical ontology's.

⁴³ <http://www.estrategy.gov/> "Extensible Markup Language (XML) embodies the potential to alleviate many of the interoperability problems associated with the sharing of documents and data. Realizing the potential requires cooperation not only within but also across organizations. Our purpose is to facilitate the efficient and effective use of XML through cooperative efforts among government agencies, including partnerships with commercial and industrial organizations."

XML / Web Services Enablement Conceptual Market Map



Copyright © 2002, ZapThink, LLC

Figure 25 ⁴⁴

Under vertical markets one would locate such well accepted industry standards such as the Association for Cooperative Operations Research and Development (ACORD), electronic data interchange (EDI), Health Level 7 (HL7) and Society for Worldwide Interbank Financial Telecommunication (SWIFT)

Professor Ben Grosoph, in his IAP course on the Semantic Web,⁴⁵ advises that the Semantic Web security and billing standards are still in flux. An estimated timeline for the development of elements making up the "Semantic Web" are identified in the figure below:

⁴⁴ <http://www.zapthink.com/reports/marketmap.html>

⁴⁵ 15.972 "Frontiers of E-Business: Introduction to Semantic Web and Web Services" in IAP 2003 (Independent Activities Period, in January).




Figure 26

Despite the immaturity of these higher-layer standards in native XML formats, it is important to recognize the changes that are occurring around us. At the RDF layer of XML that is under development, Amazon and ePinions⁴⁶, have already used RDF to link their sites to link book reviews that are written on ePinions so that they show up on the Amazon “comments” section.

The vision of the Semantic Web—which is new and still very much under construction—consists of three main parts: *ontologies*, *markup languages*, and *inference (rule-based) engines*. The first component, ontologies, are databases of information about a particular domain that are likely to be relevant to particular intelligent agents. Furthermore, these ontologies are likely to be distributed: an ontology related to medical information is likely to be maintained by a different institution than one related to tax knowledge, probably in a different city. ... The means for communicating over the Web is provided by markup languages.

The basic idea of XML-based markup languages is to provide for ease of communication across the Internet. Perhaps there will be a one-to-one mapping between ontologies and markup languages. In any event, documents using particular markup languages are what will actually be exchanged over the Semantic Web between the ontologies and the inference engines.

So this brings us to the final main component of the Semantic Web: the inference engines. The inference engines on the Semantic Web are rule-based systems, but they involve some additional challenges. For one thing, they are reasoning over information which is coming in from all over the Web, and they won't have control over when that information arrives, how often it is updated, or even how accurate or reliable it is. The engines will therefore have to have built in mechanisms for dealing with these real-time uncertainties. On the other hand, the standardization of the data through the markup languages should make some aspects of building these systems easier than before.

⁴⁶ <http://www.epinions.com/books?tab=1>

At a very high level, then, the relationship between these three components of the Semantic Web is roughly described by the following diagram:




Figure 26

The ontologies (for this medical insurance based simple example) are based in a large number of different cities. As a result, they are updated asynchronously and under the control of different organizations. The markup languages are used to communicate between the inference engines and the ontologies. And the inference engines themselves (which also may be distributed across platforms, cities, and organizations) do most of the actual "work". Note that this diagram helps to explain why things must be structured in this manner. The medical ontology must be maintained by experts in the medical field; likewise the insurance ontology must be maintained by insurance folks. Yet to process a claim, one needs to know a little about both medicine (what precisely is the claim for) and insurance (what are the rules of this particular policy). Therein lies the challenge of the Semantic Web.⁴⁷

We can see from the subsequent World Wide Web Consortium illustration how ontology's are developed and incorporated into annotations that are in turn picked up by the inference engine and routed to the appropriate destination:

⁴⁷ http://www.ramalila.net/Adventures/AI/the_semantic_web.html




Figure 27 ⁴⁸⁴⁹

The figure shows a community portal in relation to other parts of the semantic web. As David Ewalt writes in a recent Information Week article, "The Semantic Web, The Next Step In The Web's Evolution," promises even more dramatic changes where the Semantic Web changes the approach businesses take to data in that its value would be defined less by the program it's held in. "[55] As Eric Miller of the World Wide Web consortium says "Try to realize that your data is far more important than the application that accesses it."

⁴⁸ <http://www.semanticweb.org/about.html#bigpicture>

⁴⁹ **"Semantic Web: A Certain Progress, Many Challenges" Speaker: Tim Berners-Lee; MIT LCS Abstract:** The Semantic Web as a global web of data expressed as a graph of relationships. Currently, basic data formats are maturing, ontology standards are in the works, and standards for rules and inference, and the exchange of proofs are at various stages of research and advanced development. At LCS/W3C, the RDF graph language has been extended ("Notation3") to express inference rules. Practical experience in using these for personal information management, and also in the Annotea distributed annotation system provide one set of challenges, while attempts to define a formal system which unifies all the various systems used by others provide another set. Can rules be interchanged between inference systems and database systems? ...How can one best prototype a future vast system of heterogeneous knowledge bases?

As Dennis Quan said in his presentation at the AI Lab: "Imagine a world where software will be able to automatically schedule doctor's appointments based on mutual availability, put together travel itineraries that match very specific criteria, and exchange data of various forms seamlessly--be they biological, commercial, or statistical in nature. This is what the Semantic Web--basically today's Web enhanced with machine-readable information--is trying to foster. In supplementing the human-readable content of the current Web with agent-processable metadata, the Semantic Web community will be attempting to leverage half a century of AI research on knowledge representation."⁵⁰

xBRL SEA CHANGE

This paper was prepared with the help of XML technology with an application, EndNote⁵¹ which links to libraries around the world and imports the data into the correct reference format --facilitating the task of preparing this bibliography from different library reference systems including EBSCO, Platinum et al. Adobe and now Microsoft documents all are annotated using an XML-based mark-up language. The way in which we do business will be heavily impacted by these standards.

One example came up in a recent presentation to the 15.928 Strategic Management and Consulting seminar this spring Dr. Robert Eccles, Senior Fellow of Price Waterhouse Coopers, spoke of the impact that xBRL standards could have on improving analyst reporting and the ability of companies to communicate their message effectively to the market.⁵² XBRL, the standard developed by the American Institute of Certified Public Accountants (AICPA) for financial reporting and supported by Microsoft, Morgan Stanley, Edgar Online and Reuters. XBRL is being offered as an "add-in" for Microsoft Office 11, allows analysts to import financial documents into spreadsheets to generate comparisons between firms. If the SEC were to back this standard, a true "network effect" would be created across public companies. Already the FDIC in the US and the Australian Prudential Regulatory Authority are requesting xMRL filings. A powerful demonstration of XBRL analyst tools can be found on the web at www.nasdaq.com/xbrl/ and includes automated formulas for Debt Capacity, Coverage, Profitability and Cash Flow Ratio Analysis.[56]

XML PROTOCOL STACK

⁵⁰ "Weaving Your Own Semantic Web": Dennis Quan, MIT AI Laboratory,

⁵¹ from Thompson ISI Research; www.endnote.com

⁵² May 5 2003, "Restoring Credibility and Earning Stakeholder Trust"

For practical purposes today, when we talk about making web services available it is commonly referring to the protocol stack and standard services that are depicted in the following diagram:




Figure 28 Web Services stack ref. Gartner[57]

While Web Services have universally adopted XML as the mechanism for encoding data, but an array of higher-level standards is required to make Web Services work. For example, two banks that want to communicate with one another will need more than XML. They also must agree on a higher-level standard such as OFX (Open Financial Exchange). Achieving coordination between all interested parties for high-level data standards is the single biggest difficulty Web Services faces over the near term. For this reason, we expect Web Services to be adopted within organizations first because the coordination issues can be simpler in-house than those between organizations.[58]⁵³

SECURITY

Web Services adoption hit a series of roadblocks in 2002 of which security was the most notable, which ZapThink laid out the Web Services Adoption Roadmap as seen in Figure 23. Their research was positively borne out as IT organizations realized that they had to overcome critical security and management challenges before their Web Services implementations would meet broad enterprise needs.

In early 2002, industry vendors, end-user organizations, and professional services organizations tackled the critical issues of identity management, authentication,

⁵³ Some of these **XML Standard** protocols are listed herewith: ACORD; FinXML; FIXML; FpML; IFX; IRML; ISO 15022; MarketsML; MDDL; MISMO; OFX; RETML; RIXML; STPML; SWIFTML; XBRL

authorization, administration, privacy, and confidentiality that enterprises needed before they could expose critical IT assets as Web Services. Toward the end of 2002, several vendors began roll out Web Services management products, looking to help companies remove the second roadblock to Web Services adoption... The nascent Web Services security and management solutions from 2002 will undoubtedly mature in 2003. However, security and management are just the first in a series of roadblocks to the successful adoption of Web Services and SOAs. Next on the runway: transactions, registry solutions, and Web Services orchestration and workflow solutions.

Industry groups that are focused on addressing the security issues for XML based web services include:

- Liberty Alliance: identity management solutions.
- OASIS: Service Provisioning Markup Language (SPML). The Security Standards Joint Committee promotes a unified architecture, reusable specifications and a common point of reference for initiatives supporting the secure implementation of emerging Web services.
- W3C: strong involvement in XML and related security standards.
- WfMC:TC00-1003 Workflow Reference Model.
-

The Organization for the Advancement of Structured Information Standards (OASIS) and the World Wide Web Consortium (W3C) have pinpointed 13 specifications likely to form the foundation of a Web services security framework:

A standards mishmash

The Organization for the Advancement of Structured Information Standards (OASIS) and World Wide Web Consortium (W3C) have pinpointed 13 specifications likely to form the foundation of a Web services security framework.

Standard (proposed or final)	Standards body/Status	Description
------------------------------	-----------------------	-------------

SAML Security Assertion Markup Language	OASIS/Final vote	Exchanges user and machine authentication and authorization information.
XACML Extensible Access Control Markup Language	OASIS/Committee review	Expresses policies for information access.
SPML Service Provisioning Markup Language	OASIS/In committee	Defines how to exchange user, resource and service provisioning information.
WS-Security	OASIS/Forming technical committee	Extends Simple Object Access Protocol with XML security protocols.
XrML Extensible Rights Management Language	OASIS/Gathering requirements	Manages copyrights of digital content.
XCBF XML Common Biometric Format	OASIS/Defining scope of work	Defines XML codings for Common Biometric Exchange File Format.
XML Digital Signature	W3C/Completed	Provides integrity, signature assurance and non-repudiation.
XML Encryption	W3C/Final vote in committee	Encrypts and decrypts digital content.
XKMS XML Key Management Specification	W3C/Working draft	Provides a method for obtaining cryptographic keys.
Transport Layer Security/ Secure Sockets Layer	IETF/RFC 2246	Secures Internet traffic between two points.
SASL Simple Authentication and Security Layer	IETF/RFC 2222	Adds authentication to connection-based protocols.
Kerberos	IETF/RFC 1510	Provides tickets for authenticating users.
BEEP Blocks Extensible Exchange Protocol	IETF/RFC 3080	Helps establish quality of service over the Internet.

Figure 29 ⁵⁴

We are fortunate to be able to see the fruits from some of these standards efforts emerging in conjunction with the writing of this thesis.

San Francisco, CA - April 15, 2003 (RSA Conference) Today for the first time many major companies will showcase Liberty Alliance-enabled products working together, making it easier for businesses to solve critical IT challenges associated with identity management and offer innovative new web services to customers, partners and employees. Liberty Alliance's public interoperability event, being held at this week's RSA 2003 conference, brings together 20 of the industry's leading hardware, software, mobile device and service companies. These companies will showcase how Liberty's Phase 1 specifications for opt-in account linking and simplified sign-on can be used today in numerous business scenarios..⁵⁵

⁵⁴


⁵⁵ <http://www.projectliberty.org/>; Companies participating in the event are AOL, Communicator Inc, Ericsson, HP, Jabber, Inc., Mycroft, NeuStar, Nokia, Novell, NTT, Ping Identity Corporation, Phaos Technology, PostX, SchlumbergerSema, Sigaba, Sun Microsystems, Symlabs, Trustgenix, Vodafone and Waveset. Each of these demonstrations shows the use of Liberty's version 1.1 specifications in existing products and services

SECURITY ISSUES IN FINANCIAL SERVICES

In the financial securities industry specifically, fund managers need secure access to real-time information to make informed investment decisions. However, in today's world of disparate systems, fund managers have to administer and log into separate accounts with each Wall Street broker or independent services provider, making it difficult and time consuming to access real-time information. Communicator Inc and NeuStar demonstrated at the Liberty Alliance interoperability event how their use of the Liberty specifications allows fund managers, brokers and service providers to federate identities, deliver and view aggregated information and access their separate accounts in a seamless, secure and controlled manner.

MANAGEMENT

Traditional management systems are still required not only to monitor the Web services layer, but to coordinate the receipt of information (for example, events and performance data) across the entire application stack, including the Web services layer.[59] As was indicated in the beginning of this thesis, management of XML sessions is a complex task. A debate in the industry exists as to whether these sessions can be managed effectively at the HTTP layer with load balancers or whether more specialized management is required. Gartner Group exposed the management layers nicely in the following illustration:



Source: Gartner Research

Figure 30

INFRASTRUCTURE REQUIREMENT

Say the words "Web services," and the three things that will come to most developers' minds are UDDI (Universal Description, Discovery and Integration), electronic business XML and SOAP. Not many people would include LDAP, yet the directory protocol might provide the key to building a solid Web services infrastructure.

LDAP's primary strengths—maturity, reliability, scalability and security—are still scarce in current Web services standards. These key attributes should make LDAP a major contributor in the near future, and several heavy hitters are working on doing just that.

On the standards front, Novell Inc. has issued the draft specification "LDAP Schema for UDDI," which aims to formalize the role of LDAP in Web services by extending LDAP's schema.⁵⁶In particular, the draft defines schema elements to represent the

⁵⁶ (Go to www.eweek.com/links to see a copy of the draft.)

following UDDI data types: BusinessEntity, BusinessService, BindingTemplate, tModel and Publisher Assertion.

On the product front, Sun Microsystems Inc.'s new Sun ONE (Open Net Environment) Registry Server, which made its debut early this summer in the Sun ONE Developer Platform, allows IT managers to store UDDI repositories in an LDAP directory.

In addition to its evolving function in UDDI, LDAP will continue to play a central role in Web services security by acting as a central user identity store for large-scale Web services deployments.

Meanwhile, LDAP may move toward using a Web services-based interface as a directory synchronization technology (something Novell is already doing with its DirXML) and as an alternate query mechanism to the LDAP native query APIs.

LDAP will also play a role in efforts (principally at Microsoft, IBM and Systinet Corp.) to allow Web services servers to pass user credentials among themselves to create a chain of Web services.[60]

The Extensible Markup Language (XML) protocol is revolutionizing the way applications and devices share information, but the biggest change may be felt in the infrastructure. With their increased bulk and complexity, XML messages can bring application servers to a crawl, load balancers to a stop, and run data stores aground. While I do not expect to convince all of my colleagues at MIT who are "Fat Dumb Pipe"⁵⁷ advocates of the necessity for this management layer, I do suggest that any group responsible for this traffic put in place the appropriate test methodologies so as to be able to make informed decisions.

TESTING

To evaluate the performance of Web services platforms, several benchmarks are useful:

1. Number of messages per second successfully processed
2. Consistent processing of messages under increased load
3. Differences in performance between processing of the DOC/LIT and RPC/ENC encoding models

The number of messages per second a product can process is dependent on the XML parser employed. The SOAP request must be decoded by the parser and then handed off to the appropriate Java method. An inefficient XML parser will degrade

⁵⁷ Professor David Clarke would be pleased with the ministrations of my MOT colleague Richard Willey in this regard.

performance significantly as the size of the request increases, because there is more XML to decode.

Also entering the equation is the way in which a platform marshals and unmarshals arguments--in English, the way a platform transforms an XML element into a Java argument and vice versa. The way in which this is handled varies according to implementation. For example, Cape Clear's implementation handles DOC/LIT differently from the way it does RPC/ENC. In the DOC/LIT model, every argument--even primitive data types, such as an integer--is encoded into a corresponding Java class. That class may be simple, consisting of a single piece of data, or complex, a complete object that may comprise individual objects within it. In the RPC/ENC model, the arguments are often passed to the Web service as a primitive, which can increase performance.

Some platforms, such as those from Novell or Systinet, use the DOC/LIT model of processing even for RPC/ENC, which can degrade performance because of the additional instantiation of objects that are, in some regard, spurious. BEA provides an extra level of abstraction that allows elements within a SOAP request to be mapped to Java methods or arguments that do not necessarily share the same name. This is a definite value-add, making BEA's implementation more flexible, but it does add additional time to the process of marshaling and unmarshaling arguments.

Take advantage of the tuning suggestions made by your vendor of choice. Changes to the Java Virtual Machine parameters, such as memory or threads, and modifications of activation policies, log levels and connection-level settings can drastically increase the performance of an application server, so consider tuning part of the deployment process[53]

If in fact you decide to secure and manage XML sessions, two scenarios are being developed to address these issues, one based on XML proxies running on specialized hardware and the other based in software.

6. XML PROXIES

HARDWARE XML SERVICES

Our initial concern about the ability of network infrastructure to support XML sessions completion is being addressed by a variety of hardware vendors today. For example an immediate hurdle we can identify is that today's Cisco and F5 switch performance crumbles when high density SSL processing is added to the switching overhead.

Validation for the overall problem came in to the MIT \$50K Semantic Switch project from a variety of sources, most recently from Milan Jelaka, an Oracle analyst, who wrote in his communication:

“While reading the document I thought that your solution could actually complement that of Oracle quite nicely, if you manage to deliver fast and secure connection between two parties involved in exchanging transactions over the web. Speed is critical, especially if we talk about real time response. Recently I had this case at one of my client's where they wanted to establish fast and secure connection between its enterprise application and that of its partner. Namely, my client out-sources inventory management, manufacturing and shipping to its partner, but they have problems controlling the business processes. They do not have real time information of what is available in the warehouse, what is being manufactured and what is being shipped. Unfortunately, we had neither time nor resources to develop full-blown web services solution. Instead, we deliver a half-solution that enables them to have some visibility, although not in real time. Basically the main problem was speed of communication channel between two systems. “Had there been some kind of technology which would enable fast and reliable connection and information exchange in XML format between two systems, the better solution would have been possible.”⁵⁸

The issues that customers face in deploying web services based on XML standards are various and include:

- Load balancing- to intelligently manage and distribute traffic among servers
- Cascading – to group multiple units together for overflow
- Clustering and failover- to redistribute load in case of failure
- Error recovery – to detect outages and redirect queries before user notification

The way in which to solve these performance problems is through the use of XML proxies, either in software or hardware, to process the metadata associated with these sessions.

⁵⁸ Milan Jaka, Oracle Analyst

ZapThink also saw the emergence of XML Proxies, a new category of product early in 2002 targeted at addressing one of the primary drawbacks of XML and Web Services: their inefficiency on the network. One of the greatest hidden dangers of widespread XML and Web Services usage will be their impact on corporate networks, from bandwidth as well as processor and storage perspectives. Indeed, ZapThink found that up to 25% of corporate bandwidth will be consumed by XML traffic by 2006. With the emergence of both hardware and software solutions XML proxies, IT architects and data center administrators have another weapon in their arsenal for reducing the impact of XML and Web Services on the network. In 2002 there were but a handful of vendors pursuing an XML Proxy product strategy, but ZapThink expects this segment of the market to significantly heat up in 2003⁵⁹

As indicated in the chart below web services will require network infrastructure capable of significant parsing, streaming of long SOAP message formats, and applications level intelligence for processing metadata. Seven start-ups claim to have the answer (see table): Their value proposition is based on moving XML processing off the application and Web servers in order to improve performance and save valuable processing cycles.

⁵⁹ . <http://www.zapthink.com/flash/07012002Flash.html> ZapFlash Note: Leading EDI VAN vendors include GE Global eXchange Services (GXS), Get2Connect.net, EasyLink Services, IDDX and Sterling Commerce. In recent years, however, technical innovation and increasing Internet ubiquity have fostered new ways of conducting old business. EDI VANs are now in direct competition with Web services networks (WSNs) and transaction delivery networks (TDNs). Unlike EDI VANs, which most often are private networks, WSNs and TDNs promise transparent, deterministic application integration using the opaque, indeterminate Internet. TDN vendors include Bang Networks, bTrade, Datawire Communication Networks, Slam Dunk Networks and TransactPlus .. WSN vendors include Blue Titan Software, Flamenco Networks, Grand Central Communications and Primordial

61. B. Lheureux, F.K., *EDI to Web Services Networks: Contrasts in Convergence* Transaction delivery networks and Web services networks may look different from traditional electronic data interchange value-added networks; however, their overall feature set is deceptively similar and consolidating rapidly. 24 July 2002, Gartner Report.

A Sampling of XML Accessories						
Vendor	Product / Ship date	Type	Schema Validation	XML Security	XML Switching	Transformation
Cisco www.cisco.com	CTE 1400 current	Appliance	Y	N	N	Y
DataPower www.datapower.com	XA35 XML Accelerator current	Appliance	Y	N	Y	Y
F5 Networks www.f5networks.com	BIG-IP series current	Appliance	N	N	Y	N
Forum Systems www.forum-systems.com	Sentry 1500 current	Appliance	Y	Y	Y	Y
Quadrasis www.quadrasis.com	SOAP Content Inspector (SCI) current	Software	Y	Y	Y	N
Sarvega www.sarvega.com	XP2000 current	Appliance	Y	Y	N	Y
Tarari www.tarari.com	Tarari Content Processors March	Subsystem	Y	Y	Y	Y
Westbridge Technology www.westbridgetech.com	XML Message Server current	Software	Y	Y	Y	N

Accessorize it. Most XML accessories target security, schema validation, transformation acceleration and XML switching.

Figure 31 Data Communications

As one analyst indicates: These vendors must be onto something because major infrastructure suppliers such as F5 Networks (www.f5.com) and Cisco Systems are either releasing comparable products or adjusting their marketing strategies to target the same market with existing products.

As for the network managers who'll have to help the application folk sort through this XML quagmire, the solution starts with boning up on why corporate types are going gaga over XML. The protocol promises tremendous flexibility, but at a price. Large XML messages stress the network, application servers, and data stores that they run through, and can pose security concerns as well. The new network appliances aim to alleviate these stress points. [62]

In an effort to specify infrastructure requirements in large IP networks for the support of XML transactions, I took a group of MIT colleagues to meet with Mike Afergan, architect of Akamai web services on September 10 of last year to explore the need for such a service.⁶⁰ As this thesis is coming to it's conclusion, Akamai has announced a Web Services support plan in partnership with IBM as reported in the Boston Globe.⁶¹[63].

⁶⁰ September.10, 12PM MIT Kendall Square café with fellow MOT classmates Richard Willey and Lon Wu; NOTE: Akamai President Paul Sagan, who I met in Tom Eisenman's class at HBS, has responded that they had discussed an XML service prior to my encounter with Mike Afergan.

⁶¹ http://www.boston.com/dailyglobe2/121/business/Akamai_IBM_to_launch_service+.shtml

SOFTWARE XML SERVICES

As has been suggested in chapter 4, a plethora of web services software applications have emerged which can be visualized as a series of small “S” curves building on one another, which each consolidate under their own competitive pressures.

An example of one such small “S” curve was the emergence and consolidation of vendors in the XML Data Store market. At the tail end of 2001, the Native XML Data Store (NXD) market had emerged as a separate segment, but it became clear by the end of 2002 that this segment was short-lived. As RDBMS and Content Management vendors realized the value of adding native XML data storage capabilities to their products, the NXD market segment soon became subsumed within these other markets.

Further evidence of this consolidation and evolution includes the Progress Software acquisition of Excelon, the Interwoven acquisition of XYZFind as well as the addition of native XML indexing by Sybase, soon followed by Microsoft and IBM. This trend can be viewed as similar to the commoditization of web site design by consulting firms including Razorfish, Sapient.

As software solution to the XML processing overload can run in clustered server environments so as to allow stream-based XML processing and parsing without large cached memory document files, the processing of large SOAP messages via threaded resources and cached connections. The algorithms can be run of general-purpose processors, as is the case for Google applications.

Our specific focus for the financial services On Line Transaction Processing (OLTP) industry is centered on Enterprise Integration Applications (EIA) and/or integration tool sets for legacy environments. The Labs of Network Computing has evaluated web services EIA development environments and rated them in the following table.

REPORT CARD		Web Services				
	Novell Extend Application Server 4.0	BEA Systems WebLogic 7.0 with WebLogic Workshop	Cape Clear Software Cape Clear 4	Systinet WASP Server for Java 4.5 and WASP Server for C++ 4.5	Sun Microsystems ONE Application Server 7	Iona Orbix E2A Web Services Integration Platform, XMLBus 5.4
FEATURES (40%)						
MANAGEMENT (10%)	5	5	3	4	3	4
MONITORING (10%)	4	4	2	4	2	3
AVAILABILITY (5%)	5	5	3	0	0	0
DEPLOYMENT (5%)	4	5	5	4	3	4
DEVELOPMENT (5%)	4	4	5	2	3	2
MESSAGING/DATABASE SUPPORT (5%)	4	4	3	2	4	2
INTEROPERABILITY (20%)						
CLIENT (10%)	4	5	5	5	4	5
STANDARDS (10%)	4	4	5	4	4	4
SECURITY (15%)						
AUTHENTICATION/AUTHORIZATION (5%)	4	4	4	4	4	2
SSL (5%)	5	5	5	5	5	5
XML (5%)	1	0	2	2	0	0
PERFORMANCE (10%)						
DOC/LIT (5%)	3	2	3	4	5	2
RPC/ENC (5%)	4	2	3	3	5	2
PRICE (10%)						
PRICE DEPLOYMENT (5%)	2	2	4	3	4	3
PRICE DEVELOPMENT (5%)	4	2	3	5	5	3
PLATFORM SUPPORT (5%)	4	4	4	5	4	4
TOTAL SCORE (100%)	3.90	3.75	3.70	3.65	3.40	3.05
<small>A>4.3, B>3.5, C>2.5, D>1.5, F<1.5 A-C GRADES INCLUDE + OR - IN THEIR RANGES. TOTAL SCORES AND WEIGHTED SCORES ARE BASED ON A SCALE OF 0-5.</small>						
<div style="display: flex; justify-content: space-around; align-items: center;"> B B B B- C+ C+ </div>						
<p>FEATURES based on six areas: Management reflects ease of use, consolidated management of multiple servers and the configuration options possible from console. Monitoring reflects monitoring of Web service-specific data; for a perfect score, the product should go beyond endpoint logging. Availability reflects support of load balancing, clustering and failover. Development reflects ease of use, supported methods of building Web services, generation of test clients, environment and availability of plug-ins to other IDEs. Messaging/database support is required for JMS, MQ Series, DataSources and direct JDBC connectivity to multiple databases.</p> <p>INTEROPERABILITY support of Web services standards (WSDL, UDDI, SOAP) and the interoperability of deployed services with .Net and J2EE clients from multiple platforms.</p> <p>SECURITY Authorization/authentication refers to support for LDAP, ADS, IP-based ACLs. Message-level security at the server level is a must for a perfect score. A perfect XML score requires support for XML-SIG/XML-Encryption at the application-server level. Support for SSL, both client and server certificates, required for a perfect score.</p>						
<small>Customize the results of this report card using the Interactive Report Card®, a Java applet at www.nwc.com.</small>						

Figure 32 ⁶²

While these EIA environments each have their strengths and weaknesses, all of them rely on the model of building an object layer over pre-existing maize of legacy software. As will be seen in Chapter 8, a radically architectural simplification is proposed in chapter 8 to build an XML SOAP service as an underlying CICS service.

NEW WEB SERVICES BUSINESSES

⁶² <http://img.cmpnet.com/nc/1406/graphics/1406f2report.gif>

In addition to XML proxies, software based XML service offerings can include a variety of features. The market opportunity is the creation of a Web services network, where Web services are managed by a service provider or by a product that features extensive management facilities.

- Type 1 Web service networks are product-based and are licensed much like application servers (e.g., Talking Blocks or Interkeel).
- Type 2 Web service networks are service-based and managed as an outsourced facility for coordinating partners and the services between them (e.g., Grand Central or Flamenco Networks).

The Web services network provides a bridge between services so that interoperability is ensured. It exists at a logical level above the other services, so that it can act as a central control. The management platform handles:

- Quality-of-service guarantees
- Message delivery guarantees
- Metering of service activities for the purpose of billing
- Security management
- Registration of trading partners
- Management of supply chain relationships

This [EIA] platform, along with the producer platform, represents the best market opportunity for new Web services vendors, because no group of vendors dominates it. However, the long-term ability of Web service network vendors to deliver the management platform is dependent on the vendor differentiating products from major platform vendors (e.g., IBM, BEA), whose provider platforms will evolve to include many management features. Type 2 service-based networks will need to evolve beyond value-added networks that deliver documents and into transactional service environments before they can deliver value for complex systems.” [64]

Value Added Networks (VAN) continue to evolve from their EDI roots to address a variety of business purchasing and auction bidding models.⁶³

7. Causal Loop Modeling and Scenarios

In this chapter we take web services technology and look at its impact in the delivery of financial services through the lens of a specific network business model. The fundamental strategic insight here regarding the importance of modeling was reinforced during a visit by Mike Harris, CEO of EGG, the leading on-line financial service company in Europe, to Henry Weil's 15.928 Competitive Strategy class.⁶⁴ Mike Harris described the critical importance of establishing a business model independent of a company structure in order to analyze the effects of various market conditions on company strategy.




Figure 33⁶⁵

The fundamental insight here is that rather than jumping in to “fix” company structure it is important to recognize changes in company behavior as feedback on strategic choices which can be modeled. Therefore the subsequent modeling exercise looks at the impact of industry and technology trends on a web services network business model.

BUSINESS MODEL

⁶⁴ Mike Harris lecture, Chairman EGG Financial Services, 15.928 Competitive Strategy 5/6/03

⁶⁵ From Mike Harris lecture.

– Likely response for new business penetration

We used this model to perform sensitivity analysis across these internal and competitive dynamics variables in four different market scenarios as follows:

Old Business	New Business
↑	↑
↑	↓
↓	↑
↓	↓

Figure 35

In other words we simulated up and down markets for each market served and trialed each strategic variable, whether or not to cross subsidize from the old business to the new business, to invest in Brand or Products, etc. to determine the sensitivity of each strategy to market conditions. Our findings overall show that in a declining market it is imperative to take out costs and invest in new business areas, as long as a new market has been identified that is growing.

WEB SERVICES COMMODITIZATION

The rate of commoditization in IT services is widely recognized as identified herewith in the recent issue of Information Week: “Downturn Hits IT Services Hard.”[65] This confirms the analysis we performed whereby, in the absence of an expanding market, an attempt to move “up market” by off-shore consultancies could cause a commoditization of the higher end of consulting services.

This insight was further validated, during a visit by the CTO of General Electric, Larry Biagini,⁶⁷ who said that they were moving to a fixed rate “\$ per man day” model for outsourcing IT services. When we model these we can see commoditization disrupting the industry value chain and watch the growth of competitors who had heretofore been excluded from regional markets. The emergence of the Indian contract programming industry on Wall Street is having this disruptive impact on software programming, applications development and Outsourced Business Processing today in mature IT services, and will be subject to the

⁶⁷ Visit 5/10/2003 organized by Professor Christian Dussard.

same disruptive effects of web services technology unless they act preemptively to cannibalize their COBOL maintenance business. [66]

DISRUPTIVE APPLICATIONS

At the same time as commoditization is occurring, we can see from the simulation how new links can also be formed as businesses implement new technology in growing NB markets. Here is a new kind of competitive force at work in networks where transaction costs are so low that “weak” connections can be leveraged to create new business.

Tony D’Agostino, COO of Wachovia Securities, Inc. says “web services offer a faster connection, a less-expensive way to deliver the First Call data than the alternative- a dedicated connection requiring a leased phone line, a special purpose server, and more complicated integration.”[67]. For example, i-Deal is using Web services to provide a neutral platform for connecting broker/dealers in the financial industry. For uBid, an online auction powerhouse, the benefits have been dramatic. It now takes only two weeks to add aggregators, or auction “middlemen,” to its system. Buyers, of course, have instant access.[68]

An overview of the i-Deal “greenfield” financial services application is the subject of the following Gartner report: Broker/dealers and other financial services professionals need a way to deal with many different kinds of transactions, some of which happen in minutes, while others can take days.

Objective: i-Deal's goal was to provide a neutral platform so that broker/dealers could interact in the market for new-issue transactions pertaining to fixed income and equity. With no such neutral platform available, i-Deal needed to find a way to accommodate multiple types of users (such as sales representatives and investors) and partners (including brokerages and Microsoft). i-Deal's goal also included providing a brandable platform technology, described as “i-Deal inside,” that it could sell to its customers as a private-label application.

Bottom Line: Enterprises can use Web services technology over the public Internet now, provided appropriate attention is paid to security issues. Gartner recommends that enterprises plan upfront and employ external consultants familiar with leadingedge technology before launching enterprise-class applications that utilize Web services.[69]

While it is true that new entrants are devising new ways of serving customers with Web services technologies, one must always remember that the incumbent can react, especially when they are as big as IBM, with an alternate scenario.

COMPETITIVE DYNAMICS FROM INCUMBENTS

Lest we forget that legacy incumbents can compete effectively against new entrants, a story from classmate Luis Bettencourt, MOT 2003 illustrates IBM's mastery over extending the life of IBM 360 applications. Most recently through the introduction of ODBC interface to DB2 databases, IBM obviated the need for adjunct Oracle-based data mining solutions as users took to mining data directly from the mainframe. Another example of this was the short-lived XML storage market, which IBM crushed by providing file systems with hierarchical file management that support XML.

Professor Mark Meyer commented that IBM's core strategy for the top Fortune 400 companies is the migration strategy for the Z900 "grid computing" system which maintains binary compatibility across platforms. This solution does not resolve the underlying instability present in many legacy software environments which we address with our proposal for a CICS SOAP service to follow.⁶⁸

⁶⁸ Thesis draft comments, Mark Meyer email of 5/6/03

8. Architectural Recommendations

KEY QUESTIONS

Gartner reports that the biggest questions facing IT architects today are as follows:

- How will enterprises make middleware and architecture choices that work despite the ongoing industry transition?
- What role will message brokers, service-oriented architectures, data warehouses, data integration and other strategies play in enterprise application design?

DATABASE ADAPTER CONNECTIONS

Often, applications and databases tend to be connected by lots of adapters—software that directly connects two specific applications together. But that quickly creates an unmanageable tangle of software and traffic.

"We have 160 applications. Every application has to talk to every other application," says Dan Gallaher, information technology director for Jefferson County, Colo. To get every application talking to every other application, Gallaher figured he'd have to set up almost 13,000 adapters. (In a set of 160 applications, each application would have to connect to the 159 other applications—for 25,440 individual connections and 12,720 pairs). "We looked at this and said, 'Guys, this is going to be impossible. We're going to spend the rest of our lives doing this and maintaining it.'"

Not only would it be difficult to code and maintain that much software successfully, but having that many applications talking directly to each other would create a great deal of excess network traffic, as the same information required by multiple applications is re-broadcast from the source. If every application were calling directly to every other application, it would create a huge traffic jam

In one such client application that we are working with a single IBM mainframe terminal view requires 8 emulation screens to be opened concurrently for remote data entry to occur into each field, while each application requires 8 terminal views to be open simultaneously, i.e. 64 terminal views are required for one web account entry query.

PUBLISH AND SUBSCRIBE MESSAGING

One way to eliminate that problem is to avoid adapters altogether. The key is to strip out new data being entered into a system and broadcast it throughout a network. Any system that's been programmed to receive a certain piece of data can then tune in to the broadcast and

receive the information. Feld says, "Just hang it out there, and whatever has authorization can tune into that. It doesn't jam up the bandwidth sending it out to a million people."

"Publish and subscribe messaging," works like pay-per-view cable: applications that require information broadcast by a data source "subscribe" to that source; middleware then watches network traffic for messages from sources to which an application is subscribed, and alerts the application when a message is received. Messaging middleware from Tibco, and the Java Messaging Service are typical choices for such applications.[70]

The publish and subscribe web services model therefore seems more effective than a TP based model. Further investigation with the customer will be required to determine whether, for example, context sensitive data entry requires this additional service layer.

BUILD & MAINTAIN A PLATFORM INDEPENDENT SERVICE MODEL

In the traditional "design-build-run" waterfall software lifecycle, architects create models during the design time phase for developers to follow during the runtime phase, and the models are then put on a shelf. In an SOA, however, the model continues to serve a critical role on an ongoing basis. The core model in an SOA models the coarse-grained business Services that encapsulate and expose IT functionality to the business realm. This model, which ZapThink calls the Service model, acts as the clearinghouse for information about what's going on in the IT environment at any point in time. It tracks current and future business requirements, and follows the current and future versions of the Services available on the network.[71]

In accord with recommendations from the Gartner Group and ZapThink, the author recognizes the importance of encapsulating business processes and exposing them as a web service. How to structure the access to these services becomes the chief architectural concern. To look at this concern effectively, it is useful to revisit the basic components of a service-oriented architecture (SOA).

SOA components leverage two important design principles:

- Modularity — dividing big problems into smaller problems
- Encapsulation — hiding the data and logic in each module (or "black box") from uncontrolled external access

Client and server modules communicate with each other only through documented interface “contracts.”

Encapsulation helps protect data and code in the components from careless or malicious misuse. It also shields external developers from having to understand the internals of a module they are leveraging.

SOA clarifies design, makes it easier for large or geographically distributed development teams to work together, and enables software reuse because modules can be reconfigured and tailored in new ways to meet new purposes.

SOA uses some of the same principles as object-oriented (OO) programming, but does not use the OO concept of implementation inheritance, so it is not the same as OO. SOA can be implemented with OO programs or non-OO programs.[72]

Typical benefits include:

- Reduced operational costs, through saved labor hours in manual data capture and input
- More-accurate and efficient data maintenance
- Expected application of Web services technologies to other business areas (such as funds and banking) to create revenue opportunities by linking its financial products and services
- More-timely and accurate communication with pension plan clients
- Greater visibility within the industry, leading to opportunities for other financial services companies and distributors to work with web services applications[73]

THE CASE FOR SERVICE ORIENTED ARCHITECTURE

Gartner and ZapThink research thrust is that leading-edge IS departments across industries will move to SOA and Web services in 2003. Most new releases of software infrastructure products and packaged applications will adopt or upgrade their support of SOA and Web services. Concurrently Web services will increasingly emerge as a low-cost, near-commodity technology embedded in virtually every software product. SOA adoption will accelerate, and its benefits, best practices and issues will be better understood by a growing number of users, systems integrators and software vendors.

THE CASE FOR ‘PRIMITIVES’ BASED STANDARD INTERFACES

Many of the CTO, Chief Architect, Vice Presidents of eBusiness and J2EE engineers I have met with fine the discussion around Service Oriented Architectures and dominant platforms (Websphere et. al.) to be peripheral to solving core business problems at hand. By Definition the interfaces are standard. Why not get to the legacy system primitives and create the web services from there?

In so far as transactions are concerned, native J2EE is clunky and does not scale well. Major project implementations often involve significant customization of web services software and require coding of the basic transport to optimize performance. Two recent examples are the J2EE infrastructure for the New York and Chicago Mercantile Exchanges,⁶⁹ as well as Access Systems J2EE back end infrastructure for NTT Docomo.⁷⁰ The key here is that Web services are pivotal to the future of actively managed, flexible and responsive enterprise architecture. It is not possible to sit back while competitors explore this territory without the risk that they will open a strategic business lead. This is because moving to an SOA approach is not simply a matter of spending money on new software. It is as much a model, culture and philosophy change for a software development or application deployment organization. All staff, including the management, must learn how to get the most out of this new approach. [74]

ARCHITECTURAL OPPORTUNITIES:

In Clayton Christiansen's latest work,⁷¹ he suggests that architectural initiatives with lower cost and functionality but which build on existing infrastructure are likeliest to succeed. Furthermore we recognize in the work of Utter back, Christiansen and Henderson that even slight changes in architecture (with the exception of the transition from VCR to the DVD where key players remained the same) often create opportunities for new firms to become dominant players.

What is the impact on not only the IT model, but also on the culture and philosophy of the network business? In Gartner's recommendations for adopting a Standards Based XML Interface scenario, IT organizations would be best served by hiring, whether on contract basis or permanently, advanced JAVA engineers to implement highly performance solutions. Because the standards are open and best practices require adequate code documentation,

⁶⁹ Architected by Raefer Gabriel, CTO of TrueExchange

⁷⁰ As architected by Othman Laraki of Access Systems

⁷¹ As presented at the 2003 Harvard Business School Business Pan Competition and Entrepreneurship Conference Saturday, April 5, 2003

users have the opportunity to contract the construction a low cost state of the art web services infrastructure. By contrast a packaged suite adoption is often best implemented by outside consultants who specialize in specific application suites.

ISV's LEAD THE WAY

The most popular approaches to exposing Web services from mainframe applications are being led by the independent software vendor (ISV) community. The basic architectural idea being put forth in this case is to develop web services as a CICS application running SOAP over HTTP, resulting in radical simplification and cost savings; the application would perform following functions:

- Perform the serialization/de-serialization process
- Resolve ASCII to EBCDIC coding differences
- Create a SOAP-handling capability in CICS
- SOAP handler to isolate all marshalling/unmarshalling code and routing functions

from the existing programs

IBM SOLUTION 3 YEARS OUT

In a research Note Gartner asks this question: "Web Services and the Mainframe: Reality or Ridiculous? Web services represent a new paradigm for connecting business processes independent of their underlying implementation. This approach can also be used to extend the life of legacy systems." [75]

IBM has added new features in COBOL (XML), a high-performance, nonvalidating mainframe-resident XML parsing capability, a CICS 3270 Bridge for accessing existing transactions, IMS OpenTransaction Manager Access (OTMA), and template-based XML generation (EXEC CICS DOCUMENT). However, IBM does not currently provide SOAP-handling capability directly in CICS or IMS. While we expect this capability will be included in future versions of CICS, it is realistically at least two to three years away, and IS organizations need not wait.^{72, 73}

⁷² Many of the current development tools focus on automating the generation of the necessary XML structures necessary to support this technology. Also required is a SOAP-listening or SOAP handler runtime capability that can serialize/de-serialize XML messages and route processing to the correct processing logic

⁷³ The Microsoft .NET runtime or any Java application server can provide this capability. However, on a mainframe, the problem changes: 1) The serialization/de-serialization process must resolve ASCII to EBCDIC coding differences, and 2) the mainframe does not yet have SOAP-handling capability. Implementation of Web services from mainframe functions therefore requires either:

- An adapter/connector technology to interact with the established system, or
- Creation of a SOAP-handling capability CICS, the leading transaction monitor in IBM mainframe environments, does not provide this capability, nor does IMS. IS organizations could use WebSphere Application Server on the mainframe and connect existing CICS or IMS systems through CICS Transaction Gateway (EXCI) or IMS Connect as the

While it is not impossible that IBM would reverse course, current strategy seems to be to market Websphere as the platform of choice for XML services.

END CUSTOMER WEB SERVICES TRADE-OFFS

We need two things to implement Web services: an environment in which to develop them, and a platform on which to deploy them. Choosing can be tricky--it's a lot more complicated than deciding between Sun Microsystems' J2EE (Java 2, Enterprise Edition) and Microsoft's .Net. On one side of the table are extensions to existing application servers from Apache, BEA Systems, IBM, Novell and Sun. On the other side are the platforms--products from Cape Clear Software, Iona Technologies and Systinet Corp. designed to stand on their own or reside on an existing J2EE-compliant application server. And sitting in the middle is Microsoft with its own flavor of Web services built into Windows Server 200

connector technology.

9. Case Study – ISV CICS Web Service Proposal

Enterprises seeking a Web services vendor must either 1) select a large vendor (as in Figure 2) that can provide a wide range of platform support, or 2) choose a specialized vendor (e.g., SilverStream, Grand Central, Bowstreet, Asera or Avinon) for more point-oriented approaches to one of the platforms. Repeating from an earlier citation, Gartner analyst Roy Schulte predicted that low-cost, Web services-native integration products would cut into the business of more comprehensive middleware products from companies like See Beyond, Tibco, and Vitria, because many projects do not demand the full range of features...⁷⁴ In this chapter we set out to define the business requirements, and a systems architecture which meets those requirements.

INTEGRATE ACROSS WHATEVER IS NOT GOOD ENOUGH

One critical decision companies face when creating an innovation-driven growth business is determining its optimal scope. Specifically, which activities need to be managed internally and which can be safely outsourced? ...Even in a modular architecture, successful companies still are integrated—just in a different place.

In general, companies must be integrated across whatever interface drives performance along the dimension that customers value. In an industry's early days, integration typically needs to occur across interfaces that drive raw performance—for example, design and assembly. Once a product's basic performance is more than good enough, competition forces firms to compete on convenience or customization. In these situations, specialist firms emerge and the necessary locus of integration typically shifts to the interface with the customer.

BUSINESS OBJECTIVES

Typical objectives for web services implementations at financial services institutions include:

- Improving efficiencies internally and with its clients by automating the business processes underlying pension plan updates and changes to lower labor and operating costs

⁷⁴ http://news.com.com/2100-1001-978608.html?tag=fd_top

- Linking key products and services together (for example, among life insurance policies, pension funds and banking)
- Developing a better way to offer its products to clients and distributors, such as independent financial advisors

In order to meet these objectives, during the course of working on this thesis I came up with the basic product concept to develop web services as a CICS application, resulting in radical simplification and cost savings.

CICS WEB SERVICES ARCHITECTURE

What is “not good enough” in many highly customized mainframe environments is that core processes and services are dependent on unstable (and unknown) applications layers. An “Open Systems” CICS SOAP Service application tackles the heart of the issue by writing the XML interface at the source. This specific application would perform the following functions as per the following outline and diagrams below:

- Performs the serialization/de-serialization process
- Resolves ASCII to EBCDIC coding differences
- Creates a SOAP-handling capability in CICS
- SOAP handler to isolate all marshalling/unmarshalling code and routing functions from the existing programs
- Utilizes CICS Web Support (CWS) to handle the HTTP communication (required); These CICS transactions may be 3270/BMS applications or COMMAREA transactions
- Or supports 3270/BMS transactions through the CICS TS 3270 bridge; utilizing IBM's MQSeries to communicate to the middle tier
- Includes ability to refine the granularity of established systems in COBOL prior to exposing them as Web services

Traditional Web Services Application Suite:




Figure 36

Proposed CICS Web Service




Figure 37

END USER IMPLEMENTATION STRATEGY

An implementation plan would consist of the following elements:

- Create a detailed design and prototype
- Build and test units and integrate systems
- Define ROI and develop the business case
- Deploy the system and train users

The economic rationale for developing a CICS Web Service and including development costs in Capital budget versus licensing an EIA package with annual license fees is documented in the following BASELINE Project Planners from Gartner. This costs basis must be why so many mainframe shops are running customized software like “BOSS” today.

In the example provided in Appendix 2, a \$630,000 EIA annual license fee expense is incurred, which must be multiplied by the number of systems. By contrast development costs can either be expensed in the first year or depreciated over 5 years per Gartner.⁷⁵

Many organizations are faced with the need to reduce costs while there is an urgent requirement to bring legacy systems information onto corporate networks. Adding basic SOAP and WSDL support to a product may seem relatively easy, and may seem sufficient for customer integration needs. Unfortunately neither is true. Delivering a useable, commercial-grade Web services stack requires significant investment above and beyond rudimentary SOAP and WSDL support. In order for Web services to work, they must seamlessly interoperate with all other SOAP implementations in the enterprise, a non-trivial task that presents challenges even to some independent vendors of Web services solutions. They must also support multiple transport mechanisms and integrate with existing IT infrastructure such as security systems, application servers, and management systems. Building these features requires developers, quality assurance and support staff proficient in XML, SOAP, WSDL, Java™ 2 Platform Edition (J2EE), Microsoft®.NET, and a wide variety of other Web services-related technologies.

SOAP, WSDL and UDDI are commonly understood as the primary Web services standards that are in use today. These standards are evolving quickly, and new Web services standards are also emerging: standards for interoperability, transactions, orchestrations, security, routing and management. Some are near completion, others are in flux, and some are stalled. There are also a number of Java technologies that correspond to Web services standards: JAX-RPC, SAAJ, JAXM, JAXR, and others.

⁷⁵ Please refer to Appendix for spreadsheet project model

10. Business Strategy


This chapter is built on the competitive strategy of Michael Porter,[76]⁷⁶ C.K. Prahalad and Gary Hamel's work on Resource-Based or Core-Competency Based differentiation,[77] Mark Meyer's work on platforms[78] as well as Michael Cusumano's popular book on the subject.[79] A further evolution of this theme is exemplified in [Arnoldo C. Hax](#) and [Dean L. Li Wilde](#)'s "**The Delta Project: Discovering New Sources of Profitability in a Networked Economy**,"[12] in which the authors argue that the best way to gain competitive advantage in today's economic networks is to bond with customers rather than to beat the opposition.

Three distinct positions can all lead to outperforming a given industry over the long term:

- "Best Product" — Low cost or differentiation
- "Total Customer Solution" — Reducing customer costs or increasing their profits
- "System Lock-in" — Complementor lock-in, competitor lock-out, or proprietary standard ⁷⁷

The strategic planning process which I have adopted from Forrester Research's methodology as described in the "X Internet" report are as follows:[80]

- Identify strategic positions reflecting fundamental new sources of profitability.
- Align strategic options providing a comprehensive link between the strategic direction and execution.
- Introduce adaptive processes capable of responding continually to an uncertain environment.
- Provide metrics that are drivers for performance.



76

Porter's model of competitive dynamics highlights opportunities for high differentiation or low cost strategies.

⁷⁷ http://www.dean.com/delta_model_framework.htm

These strategic metrics can then be identified to the business entities and functions as described in this organizational framework from Gartner that follows.

		IT Strategy				
		Infrastructure	Service	Applications	Integration	Sourcing
Inputs From Business Strategy	Geographic	• Network • Dispersion	• No. of locations • Organization • Languages	• Regions • Languages • Legal	• Internal BU • External • Cross-border	• Locations • Cultures • Processes
	Governance	• BU vs. enterprise • Architecture • No. of versions	• Who decides	• Strategy • Focus • Change type	• Stovepipes • Architecture	• Strategic
	Future	• Org. plan • Architectural compliance	• Foundation	• Legacy transformation • Architecture	• Enterprise • Architecture	• Skills
	Legacy IT	• Change rate • Base cost	• Service level	• Change Rate • Maintenance	• Transform	• Internal/ External
	Virtual	• Architecture • Coordination	• Type • Levels • Cost profile	• In/out • Priority	• Architecture	• Extent • Strategy • Org. structure
	Customer	• Boundaries • What's needed	• Service level • Management	• Change input • Priority	• Client-facing • Customize	• Control
	Funding	• Operational funding	• Service level • Priority	• Change funding	• Commitment • Infrastructure	• Cost vs. value • Training • Recruitment

BU Business unit
 IT Information technology

Source: Gartner Research

Figure 4. Business Strategy Elements Mapped to IT Strategy

Figure 38 ⁷⁸[42]

Using these elements as our strategic variables we can set about planning with business units in the following areas:

□ Infrastructure operations — IT infrastructure and service bounds the infrastructure operations group. This defines how the infrastructure will look technologically and how service delivery will be managed.

□ Infrastructure applications — Application portfolio change defines how business processes will evolve, and the extent of business process integration lays out the level of complexity required. Service

In this process Gartner recommends organizations consider hiring specialized consultants for both infrastructure operations and applications. “The emerging IT environment is at a level of complexity such that efforts to build IT infrastructure and integrated applications require specialized expertise that is often unavailable internally. Good operations managers will ensure that their organizations have the ability to work effectively with and integrate the value from networks of service providers with a variety of special capabilities. In every case a significant number of in-house IT professionals and users should be included in integrated IT initiatives in order to ensure enterprise objectives are met.” - Gartner

CORE PROCESS REQUIREMENTS

The business strategy input from the left side of the diagram requires identification of core business processes and their importance in order of priority. The definition of key supporting concepts in process including: orchestration, choreography, composition, collaboration, coordination, workflow, and transactions are key to this transition. According to ZapThink the standards landscape will converge on a single choreography, orchestration, and process flow specification in the next 12-18 months.⁷⁹

For a more generalized perspective John Hagel, Chairman of EGG, the largest on-line financial services provider in the United Kingdom, recommends the division of organizations into three core business processes, one rooted in Customer Relationship Management, One in Supply Chain Management, and one core process for Product Innovation and Development[81]

The capability of delivering a web services transaction capability to new intermediary portals, and in turn to link up to third party legacy environments seems like a clear win for all participants in the revised value chain. Whether to adopt a packed solution and customize it or to go with open source products and write links to standardized interfaces which optimize

⁷⁹ These include the identification and description of key specifications in support of process design and execution concepts from standards groups including: Business Process Execution Language for Web Services (BPEL4WS), Business Process Modeling Language (BPML), Business Process Modeling Notation (BPMN), Business Transaction Protocol (BTP), Conversation Support for Web Services (CS-WS), ebXML Business Process Specification Schema (BPSS), RosettaNet Partner Interface Processes (PIPs), SOAP Conversations, Web Services Choreography Interface (WSCI), Web Services Flow Language (WSFL), WS-Coordination, WS-Transaction, WS-Reliability, WS-ReliableMessaging, WS-Addressing, Extensible Process Definition Language (XPDL), Microsoft's XLANG, and the W3C working group on Choreography.
<http://www.zapthink.com/reports/ZTR-WS108.html>

transaction speed and scalability is a decision that each MIS shop will evaluate from a risk/reward perspective. The bigger risk is that the opportunity to become the “Amazon” of specific industry sectors will be lost to other players.

As Clayton Christiansen points out in his analysis, vendors eventually tend to overshoot the needs of their customers on the high end of the performance curve, while disruptive technologies often creep up from producing lower performance products or services. In his talk at the HBS Business Plan Competition Christiansen cited the price of computing (#MIPS/square inch/\$) doubling every 18 months and expressed his sense that hardware vendors will have a hard time growing their sales. Even with the cost of disk diminishing even faster, basic cutwork economics dictate that it has become cheaper to store centrally and distribute remotely rather than to process (compress) and synchronize traffic amongst distributed nodes. The models of ASP’s who leveraged thin clients and centralized servers were successful where the Exodus model of replicating customer systems in each POP proved to be overly capital intensive.

John Landry, former CTO of Lotus and IBM, during a conversation at the MIT E-Lab CEO gathering April 9 at Sloan, cited his belief that the emergence of HDTV will drive the demand for 802.11g, DSL bandwidth exceeding that which cable and DSL vendors have installed today, and will still require the increased processing power and storage at the client to provide a smooth viewing experience of material which will be “dribbled” from a central site on a TiVo model. Since there is no driver in web services that I am aware of which would require this kind of throughput or video processing I will stick with the underlying network economics described above as they impact web services infrastructure design and applications.

PRESERVING REAL OPTIONS

The most important attribute of a go-slow implementation approach and working with existing mainframe systems on a real-options model, taking each integration step separately, is the ability to adjust development to the economic value of the application. As Don Lessard said in a recent Sloan working paper:

“The real-options approach recognizes that decisions that determine project cash flows are made sequentially over many episodes. The key insight of this approach is that uncertainty or volatility can actually increase the value of a project, as long as flexibility is preserved and

resources are not irreversibly committed. As a result, the economic value of a project when it is still relatively unformed is often greater than the discounted present value of the expected future cash flows. Value is increased by creating options for subsequent sequential choices and exercising these options in a timely fashion. Thus, sponsors seek projects that have the potential for large payoffs under particular institutional and technical circumstances.”[82]⁸⁰

In my interviews regarding web services integration within specific financial services application, I asked whether it was important to move to a Web Service SOA architecture in which each discrete business process exposes and describes itself and allows integration to a vendor’s’ de facto’ web portal and applications servers...

One of the respondents, former CTO of Hogan, answered:

“Sounds like a very academic question to me. Banking and healthcare organizations are interested in getting their business problems solved efficiently and effectively. Technology and the associated architecture is only a means to end not the end itself. The chosen architectural approach cannot be generalized, but rather must fit the problem at hand. I may not be the best person to speak with about your question. Our company is very financially successful, but we strive to find simple solutions to complex problems rather than inventing or exploiting new paradigms in system architecture.

I did serve in a role as Director of System Architecture for Hogan, but this was after the product was built. For the past several years, I have been involved in Healthcare Software that is based upon an ASP model. Gary Cohen, Chief Operating Officer; Innovative Managed Care Systems, Inc. “

Our basic recommendation in light of this comment is to ensure that business services remain the drivers for architectural choices, and not the other way around. This ASP views the choice of web services architectures to be “academic” in the context that the basic XML interface is standard, whether generated in a CICS or J2EE environment.

This core insight resonates with the thesis of John Hagel’s “Out of the Box, Strategies for Achieving Profits Today and Growth Tomorrow through Web Service,”[81] which was recommended to me by Professor Tom Eisenmann in reviewing this thesis. John uses the example of Li & Fung⁸¹ as a business which can dynamically configure service providers within a process network that has finely tuned measurement systems which can monitor performance across a variety of industries, in the case of Li & Fung, designers, fabric

⁸⁰ The real-options framework is based on the same logic as that of financial options as developed by Black and Scholes (1974). Dixit and Pindyck (1995) and Trigeorgis (1996) extend it to real options.

⁸¹ Mr.Li spoke at the MIT Sloan 50th Anniversaty Celebration last fall.

suppliers, cutters, sewers, assemblers..."To successfully exploit the economic potential available at this stage, senior managers will need to evolve to organizational forms that effectively coordinate activity across a broad range of dynamically configures service providers."⁸²

NETWORK EFFECTS

The importance of architectural choices is that they enable businesses to leverage network effects. These 'network effects' opportunities are the subject of the balance of this chapter.

"What's important in a network changes as the network scale shifts. In a network dominated by linear connectivity value growth, "content is king." That is, in such networks, there are a small number of sources (publishers or makers) of content that every user selects from. The sources compete for users based on the value of their content (published stories, published images, standardized consumer goods). Where Metcalfe's Law dominates, transactions become central. The stuff that is traded in transactions (be it email or voice mail, money, securities, contracted services, or whatnot) is king. And where the GFN law dominates, the central role is filled by jointly constructed value (such as specialized newsgroups, joint responses to RFPs, gossip, etc)."[83]

Gartner group puts it even more succinctly: "For the first time, the risk of NOT investing in XML outweighs the risk of investment."

A BUSINESS THAT VIEWS ITSELF AS A WEB SERVICE - MVNO

At the recent MIT conference, "Extreme Communications: A Radical Rethinking of Business Technology and Regulatory Strategies," I asked Lewis Norman, the CTO of Freeserve, what their competitive advantage was as a Mobile Virtual Network Operator (MVNO). His answer was that Freeserve is a web service. It was striking to find a business that defined itself according to its network architecture. In exploring what he meant by "web service", it was literally that Freeserve had installed a web services platform that integrated to legacy billing systems at Energis (a former customer whose systems I helped to plan at Ironbridge Networks). This platform, in turn, gave them the ability for their customers to form ad hoc networks as pictured in the following offer.

FreeServe Web Services Based Community Network

⁸² "Out of the Box." John Hagel, p.194.



Smartgroups.com is a community based website designed for clubs, societies, associations, or any group of friends and colleagues. It creates web files for you to store group information and allows you to keep in touch with email messaging.

Once a Smartgroup member you'll be able to access your group wherever they are, using your mobile Internet phone. You'll be able to:

- access all your groups
- search for and list group events
- view details about specific events
- search for and list messages for your group
- read emails posted to the group
- send emails to the group or other members

Figure 39 ⁸³

Here we have an example of a traditional service business; in this case, a mobile phone company who has taken the technology of web services and applied it to creating a new business model. Within 90 days of offering this service while still in beta, I understand that they have registered 90,000 user groups. These powerful network effects change the underlying economics of the mobile phone business in the same manner that we saw earlier in the MCI Friends and Family Program at a fraction of the implementation cost. As Randy Rhettberg at the MIT AI Lab pointed out in his talk on "Life at the Top of the S Curve," who would have guessed that instead of 'voice over IP' we would get pictures being transmitted over wireless MMS systems! These applications are clearly not being driven on the technical merits of the networks they happen to be running on. Despite the mismatch between infrastructure and the application, the wireless industry is leveraging network effects to drive new services revenue.⁸⁴

⁸³ <http://www.smartgroups.com/user/whyregister.cfm?>

⁸⁴ Perhaps the model being adopted by Altnet to sell licensed content and Sharman Networks key word search advertising model sponsored by Microsoft, Netflix and DirectTV offer a path to resolution⁸⁴. *Business and Money*. Boston Globe, April 11, 2003..

COLLABORATION SOFTWARE

A range of software development companies have developed products to assist companies in enabling their core business services with network capabilities including instant messaging (IM), user groups, etc. These applications include independents like the Jabber IM platform which are themselves based on distributed XML routing, and Microsoft's product roadmap which focuses on collaboration.[85] Further examples have emerged at the recent beta release of the Liberty Alliance XML Security Interoperability Forum as per the following:

White Plains, NY—April 14, 2003—Dow Jones Newswires and Communicator Inc have reached news distribution and software development agreements in which Dow Jones Newswires will deliver its financial news and information through Communicator Inc's Hub IM. The unique offering integrates Dow Jones' content with Communicator Inc's messaging and collaboration, identity management and content aggregation technologies. The result is a single, secure application through which users can receive market news and then communicate and collaborate about that information in real time.

Peer to peer transmissions, now that we are back up to the numbers of file transfers at the height of Napster, are only a shadow of what is likely to come with machine-to-machine transactions. Here in effect we may find that the "scale free" mathematical progression discovered by Albert-Laszlo Barabasi is in fact exceeded by the formula David Reed of the Media Lab puts forth:

This exponential law of GFNs, like Metcalfe's Law, creates increasing returns as scale increases, which has surprising economic results. Both laws give a powerful bonus to interconnection; mergers and partnerships of networked companies should be able to extract a premium resulting from these laws. When we combine two networks together so that users of one network can connect seamlessly to users of the other, Metcalfe's Law tells us already that substantial new value is created: $(M+N)^2 = M^2 + N^2 + 2MN$. This bonus term, $2MN$, is substantial-up to 100% of the value in the original unconnected networks. Thus there is an enormous incentive to find ways to interconnect networks, since the members of each network can access a much larger set of potential transaction partners. With the GFN law, interconnection is even more powerful, creating many new potential groups that span the two networks: $2^{M+N} = 2^M 2^N$. The GFN interconnection bonus percentage itself grows exponentially with the size of the smaller network.

What we see, then, is that there are really at least three categories of value that networks can provide: the linear value of services that are aimed at individual users, the "square" value from facilitating transactions, and exponential value from facilitating group affiliations. What's important is that *the dominant value in a typical network tends to shift from one category to another as the scale of the network increases*. Whether the growth is by incremental customer additions, or by

transparent interconnection, scale growth tends to support new categories of killer apps, and thus new competitive games.[83]

Let's recall the Internet traffic characterizations of the first chapter. It will be important to measure the impact of traffic formed by Group Forming Networks vs. batch XML machine-to-machine exchanges. What combination of dragonflies or tortoises result, be they mice -like SMS packets or elephantine-streaming video or multimedia CAD groupware?

Referring back to Charlie Fine's presentation at the "Extreme Communications: A Radical Rethinking of Business Technology and Regulatory Strategies," a framework for the creation of a value chain, in this case of web services infrastructure, involves three components:

- Insourcing/Outsourcing (the Make/Buy or Vertical Integration Decision)
- Partner Selection (choice of suppliers and partners for the chain)
- The Contractual Relationship (Arm's length, joint venture, long-term contract, strategic alliance, equity participation, etc.).

These are areas of innovation for further development, as would John Hagel's core business processes of Supply Chain Management, Customer Relationship Management and Product Development.

We have covered some of the underlying technologies that permit businesses to get close to their customer(and their customers' customers). For software infrastructure we have pinpointed the possibility of interfacing with applications at multiple levels, and the trade-off between higher-level simple user interfaces and the performance and scalability benefits of opening standard interfaces as close to the "primitives" as possible. When such an interface is published as a service there are many he benefits which revert to the firm.

From a hardware infrastructure standpoint we have covered emerging traffic trends and the new protocol sessions that XML sets up. We have asked the question about how best to gain management visibility on these sessions, from both a security and management standpoint. As an aside, the benefit of integrating within a legacy CICS environment would be the continued use of security and audit tracking based on client ID and passwords. Please note that moving access beyond the firewall presents a series of security and performance challenges, standards for which are in flux and carry a high level of risk.

From a human resources perspective, the build versus buy equation leads us to suspect that even with the expensive integration required, that application suites including EIA are likely to benefit the single mainframe user site, and by contrast that the benefits of developing a customized solution, optimized for a highly performant mature application environment, with the savings from paying for site licenses for each additional mainframe, warrant careful attention.

STRATEGY CHOICES

If every business is a network business, then every business process has the potential to become a web service which can be subscribed to. Some of the things we have learned from our experience include that it is difficult to implement new technology into legacy environments. CICS and DB2 installations are remarkably resilient. Nonetheless disruptive Greenfield competitors will emerge, as in the case of uBid, if incumbents do not get out to cannibalize legacy ways of doing business and invest in the networked world.

The IT challenge for web services is to find simple solutions to complex problems rather than inventing or exploiting new paradigms in system architecture. Ultimately, a distributed infrastructure is required to handle the speed and transaction reliability required to sustain business on the large -scale across the diversity of applications that are the focus of this study. Distributed development technologies such as DCOM, CORBA, RMI, SOAP and WSDL, essentially extensions of simple RPC protocols, provide the interconnectivity between components of such a distributed infrastructure.

In their initial instantiation in major networks today, web services can be implemented as an extension of EDI networks. The same legacy point-to-point connections and transaction networks remain undisturbed, extending web portals up and down the supply and demand chain. Back to the initial analogy with historical precedents, web services may look more like Edison's challenges of integrating with pre-existing gas distribution infrastructure rather than building a fresh set of rails. In this context we believe there is a way to integrate web services even more directly to the customer transport layer. This is a business opportunity we are pursuing in parallel with the completion of this thesis.

I hope that some of the insights from this study can be used to encourage networked businesses to evaluate using the range of web services capabilities to ensure participation in and expanding the evolving value chain of their industry.

FUTURES

One of the more interesting areas to speculate about in applying open source networked development capabilities is in the development of a global Ontology of basic information.

Cyc, which is not an acronym but rather is short for enCYClopedia, aims to capture all, or a large segment, of human knowledge such as might be found in a typical encyclopedia in a knowledge base. Cyc is the brainchild of Douglas Lenat, with significant assistance from R.V. Guha, who conceived back in 1984 of the idea of encapsulating all of human knowledge in a rule-based system. ...

Cyc is a very large knowledge base containing, in its open source version, about 6,000 concepts and 60,000 assertions about those concepts. These concepts contain basic human knowledge which is available in a typical encyclopedia. The full version of Cyc contains many more concepts, although they don't say what the precise number is. Cyc provides, in addition to the concepts, [a rule based system](#) for reasoning about the concepts (they call their rule-based system a "knowledge server"). To use Cyc, one would need to couple their general knowledge server with specifics about a particular problem to be solved. For example, let's say we know, on our own without using Cyc, that a given business is located in Chicago. Cyc will probably have concepts describing Chicago, Illinois, the Midwest, and the United States. The knowledge base will be able to infer that the business located in Chicago is also located in Illinois, the Midwest, and the U.S., without specifically having to be told those facts.⁸⁵

Applying this logic to the software business, Phil Evans points out that "scale free" network innovations from the community of Linux developers threaten Microsoft's core operating systems business. This same logic leads to the question as to whether open source web services software like Apache Tomcat undermines or complements platform leadership of the leading suites IBM Websphere, BEA WebLogic, Microsoft .NET....

At the MIT Information systems Conference "Extreme Communications: A Radical Rethinking of Business Technology and Regulatory Strategies," Frans Kaashoek of the Laboratory for Computer Sciences presented work of his colleagues⁸⁶ on a distributed lookup protocol. The contribution is a scalable protocol for look-up in a dynamic peer-to-peer

⁸⁵ <http://www.opencyc.org/> Beta Version 0.7 of OpenCyc is now available supporting Linux and Windows NT/2000/XP

⁸⁶ Ion Stoica, Robert Morris, David Liben-Nowell, David R. Karger, M. Frans Kaashoek, Frank Dabek, Hari Balakrishnan, *Chord: A Scalable Peer-to-peer Lookup Protocol for Internet Applications*. To Appear in *IEEE/ACM Transactions on Networking* [Postscript PDF](#)

system with frequent node arrivals and departures. The Chord protocol is based on a hash table architecture that is generated on a peer-to-peer basis and uses both semantic key and node information as an alternative to centralized Domain Name Server (DNS) directories.⁸⁷

The Chord project aims to build scalable, robust distributed systems using peer-to-peer ideas. The basis for much of our work is the [Chord](#) distributed hash lookup primitive. Chord is completely decentralized and symmetric, and can find data using only $\log(N)$ messages, where N is the number of nodes in the system. Chord's lookup mechanism is provably robust in the face of frequent node failures and re-joins.

One way that we use Chord is as the basis for the [CFS](#) (Cooperative File System) storage system. CFS allows anyone to publish and update their own file system, and provides read-only access to others. CFS distributes the load of serving data very widely to achieve high performance despite flash crowds. It also replicates all data, and maintains that replication as nodes fail and re-join the system.

Concluding with an observation from Randy Pond, leader of supply chain logistics at Cisco, during the MOT Silicon Valley trip: “if only that 1970’s congressman had gotten his way to have definitions of things specified in legal code, i.e. defining that ‘a pin is a pin,’ we would have set the ontological framework for vastly simplified and integrated supply chain. Rosetta.net must work!”

Meeting with Randy Scott, the CTO of General Motors was a reminder of the immensity of the task ahead. Integrating even a subset of 30,000 vendors, each of whom are currently required to implement different EDI systems for the major car OEM’s, despite the auto industry’s Connexant initiative. There is clearly a crying need to represent the next level of semantic standards. Nonetheless one can imagine the excitement that applications such as “push to talk” and Instant Messaging will have on the on a subset of GM’s 30,000 vendors as those capabilities are integrated into the Microsoft CAD package that GM is a Beta site for.

Returning to the basics of any networked business, a standard interface and ontology is the basis for suppliers and buyer exchange, be they builders of Wireless Mesh Network WiFi base stations, HDTV equipment, or participants in the latest “Match.com” on-line dating service.⁸⁸ . The further trait for that transaction to add value to the entire network, like an additional node in a mesh architecture adds value to it’s neighbors, underlies the promise of a network business model. Peer to peer and smart agent interactions represent a shadow of the promise that machine-to-machine communications represent. As industries evolve

⁸⁷ <http://p2p.semanticweb.org/>

⁸⁸ These were the examples of network businesses the students selected to present in Tom Eisenmann’s HBS class.

ontological standards for machine to machine communications, the value chains that are characteristic of Dell and Cisco will become a requirement for best practices in every industry. As was stated in the Acknowledgements section, successful societies are ones in which groups of people can form networks dynamically to align on a given purpose, and then reconfigure themselves according to revised priorities. Web services offer organizations the ability to do just that.

Bibliography Footnotes

1. Borges, J.L., *El Aleph*. 1971, Madrid,: Alianza Editorial. 184.
2. Fine, C.H., *Clockspeed : winning industry control in the age of temporary advantage*. 1998, Reading, Mass.: Perseus Books. xv, 272.
3. Munsinger, D.B. and Massachusetts Institute of Technology. Management of Technology Program., *The evolution of the wireless equipment value chain*. 2002. p. 124 leaves.
4. Kawashima, M., *Telecom Value Chain Dynamics and Carriers' Strategies in Converged Networks*, in *Management of Technology*. 2002, MIT Sloan: Cambridge MA.
5. Thomas R. Eisenmann, R.H.E., Mary Tripsas (Editor), *Internet Business Models: Texts and Cases*. 2002, New York: McGraw Hill Text.
6. Utterback, J.M., *Mastering the dynamics of innovation*. 1996, Boston, Mass.: Harvard Business School Press. xxix, 253.
7. Pine, B.J. and J.H. Gilmore, *The experience economy : work is theatre & every business a stage*. 1999, Boston: Harvard Business School Press. xii, 254.
8. Simchi-Levi, D., P. Kaminsky, and E. Simchi-Levi, *Designing and managing the supply chain : concepts, strategies, and case studies*. 2000, Boston: Irwin/McGraw-Hill. xxii, 321.
9. Rheingold, H., *Smart Mobs: The Next Social Revolution*. 2002, New York: Perseus Publishing.
10. Lisa M. Bowman Staff Writer, *Who let the NeoPets out?* February 26, 2002, CNET.COM.
11. Evans, P. and T.S. Wurster, *Blown to bits : how the new economics of information transforms strategy*. 2000, Boston, Mass.: Harvard Business School Press. xv, 261.
12. Hax, A.C. and D.L. Wilde, *The Delta project : discovering new sources of profitability in a networked economy*. 2001, New York: Palgrave. xxi, 278.
13. Barab*si, A.-L., *Linked : the new science of networks*. 2002, Cambridge, Mass.: Perseus Pub. 280.

14. Whiting, R., *Hone Business Operations with Integration Informatica and webMethods combine data integrations, EAI, and business intelligence*. Information Week, April 14 2003: p. p. 28.
15. Liebowitz, S.J., *Re-thinking the network economy : the true forces that drive the digital marketplace*. 2002, Amacom: New York. xii, 224.
16. Economist, T., *The First Will Be Last: Rethinking the Networked Economy: The Real Forces that Drive the Digital Marketplace*. The Economist, Sep 26th 2002.
17. SLIWA, C., *Q&A: GM's CTO outlines Web services strategy*. Computerworld, FEBRUARY 13, 2003.
18. Barabási, R.A.a.A.-L., *Statistical mechanics of complex networks*, in *Reviews of Modern Physics--January 2002 Volume 74, Number 1, pp. 1-294*. Published 30 January 2002.
19. Claffy, K.C., *CAIDA: VISUALIZING THE INTERNET*. Internet Computing Online, Jan 2001, Cooperative Association for Internet Data Analysis - CAIDA
San Diego Supercomputer Center,
University of California, San Diego.
20. Cheswick, H.B.B., *The Internet Mapping Project*. 2000, Lumeta Corporation.
21. Claffy, N.B.a.K., *Understanding Internet Traffic Streams: Dragonflies and Tortoises*. IEEE Communications,, October 2002. **vol. 40, num. 10**,: p. pp. 110-117.
22. Pool, I.d.S., S. Milgram, and T. Newcomb, *The Small world*. Communication and information science. 1989, Norwood, N.J.: Ablex Pub. xxvi, 382.
23. Barabási, A.-L., *The Physics of the Web; Statistical mechanics is offering new insights into the structure and dynamics of the Internet, the World Wide Web and other complex interacting systems*. PhysicsWeb, July 2001(July 2001).
24. Matlis, J., *Scale-Free Networks*. Computerworld, 2002.
25. R Albert, H.J.a.A.-L.B., *The Internet's Achilles' heel: error and attack tolerance of complex networks*. Nature, 2000. **406**: p. 378.
26. Desmond, P., *Virus inoculation: How ISPs can ease the virus burden*. 2001, eSecurity Planet.com.
27. Christensen, C.M., *The innovator's dilemma : when new technologies cause great firms to fail*. 1st HarperBusiness ed. 2000, New York: HarperBusiness. xxvii, 252.

28. Schumpeter, J.A., *Theorie der wirtschaftlichen entwicklung; eine untersuchung *ber unternehmergewinn, kapital, kredit, zins und den konjunkturzyklus*. 2. neubearb. aufl. ed. 1926, M*nchen und Leipzig,: Duncker und Humblot. xiv, 369.
29. By W. Brian Arthur, *Is the Information Revolution Dead? If history is a guide, it is not*. Business 2.0, March 2002 Issue(March 2002 Issue).
30. Allan Vermeulen (Author), S.W.A.A., Greg Bumgardner (Author), Eldon Metz (Author), Trevor Misfeldt (Author), Jim Shur (Author), Alan Vermeulen, Patrick Thompson, *The Elements of Java Style*. 2000, Cambridge, England: Cambridge University Press.
31. Fukuyama, F., *Trust : the social virtues and the creation of prosperity*. 1995, New York: Free Press. xv, 457.
32. Pender, L., *IT Versus Marketing; How to end the war and profit from a new relationship*. CIO Magazine, Jun. 15, 2001(Jun. 15, 2001).
33. Walko, J., *Motorola to buy soft switch start-up Winphoria*. CommsDesign.com, An EE Times Community, April 16, 2003.
34. Oliveau, A.M.a.-X., *Confronting the Limits of Networks*. MIT Sloan Management review, Summer 2002: p. 85-87.
35. Porter, M.E., *Competitive strategy : techniques for analyzing industries and competitors*. 1980, New York: Free Press. xx, 396.
36. Serman, J., *Business dynamics : systems thinking and modeling for a complex world*. 2000, Boston: Irwin/McGraw-Hill. xxvi, 982.
37. by Sarah Jane Johnston, H.W.K., *Publish Date: A new Harvard Business School working paper traces the evolution of IT management consulting and trends for the future. Read our e-mail interview with professor Richard Nolan and HBS Interactive Senior Vice President Larry Bennigson*. Mar 31, 2003, HBSWK.
38. Natis, Y., *The Web Services Provider Platform: You Already Have One*
A Web services provider platform hosts and provisions Web services for the enterprise, its partners and its customers.
At the technological core of the platform is service-oriented architecture. 24 April 2002, Gartner Group: Stamford CT.
39. Group, G., *Key Components for Building Your Architecture,"*. 2003, Gartner Group: Stamford, CT.

40. Ronald Schmelzer, S.A., ZapThink, LLC, *Service-Oriented Process: Thinking about Processes before Services*. April 22, 2003, ZapFlash.
41. Holst, S., *Newsletter on Content Management & Information Technologies The History of Content Management*. January 30, 2003, Gilbane.
42. Grosf, B.N., et al., *A declarative approach to business rules in contracts : courteous logic programs in XML*. eBusiness@MIT working paper ; 113. 2001, Cambridge, Mass.: MIT Sloan School of Management. 10 leaves.
43. Soat, J., *IT Architecture*. Information Week, Nov 11 2002.
44. LaMonica, M., *Flexible software architectures on rise*, in *CNETNews.com*. February 20, 2003.
45. Bloomberg, J., *Thrift: the New Normal*. February 6, 2003, ZapFlash.
46. Baldwin, C.Y. and K.B. Clark, *Design rules*. 2000, Cambridge, Mass.: MIT Press. v. <1>.
47. Pezzini, M., *Is J2EE Good Enough for OLTP?* 24 October 2001, Gartner Group: Stamford Ct.
48. Grice, B.C. and C.N.c. Staff Writer, *Microsoft committed to wireless, Gates says*. February 28, 2000, 1:30 PM PT, CNET.
49. Y. Natis, M.P., D. Plummer, R. and G.P. Schulte, C. Haight, M. Driver, *The Application Platform Suite Magic Quadrant*. 29 October 2002, Gartner Group: Stamford CT.
50. Henderson, R., *Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Industry*. Rand Journal on Economics, 1990.
51. Vecchio, D., *Programmatic Integration Servers Are a Growth Opportunity*. 20 March 2003, Gartner Group: Stamford CT.
52. J. Thompson, R.S., M. Govekar, F. Kenney, B. Lheureux, D. McCoy, Y. Natis, M. Pezzini, J. Sinur, *CIO Update: Gartner's 2Q02 Magic Quadrant for Application Integration Vendors*. 10 July 2002, Gartner Group: Stamford CT.
53. Lori MacVittie, N.C.T.E., *Serving up SOAP*, in <http://blog.networkcomputing.com/webserv.html>. April 3, 2003.
54. Madnick, S.E., *MISGUIDED SILVER BULLET: WHAT XML WILL AND WILL NOT DO TO HELP INFORMATION INTEGRATION*. Sloan Working Paper 4185-01 eBusiness@MIT, October 2001(Working Paper 111).

55. Ewalt, D.M., *The Next Web; The Semantic Web, the next step in the Web's evolution, promises even more dramatic changes*. Information Week, Oct 14, 2002.
56. Mcdougall, E.C.W.P., *Easier Financial Reporting At Hand With XBRL Nasdaq first stock market to adopt XML-based tags for financial data*. Information Week, Aug. 12, 2002.
57. Jannette Alston, D.H., Russell Ruggiero, *UDDI (Universal Description, Discovery, and Integration) and Web Services: Perspective*. 12 November 2002, Gartner Group: Stamford,CT.
58. Shirky, C., *Planning for Web Services: Obstacles and Opportunities*. May 2002, O'Reilly.
59. Cameron Haight, B.L., Whit Andrews, *Extra Effort Needed to Manage Many Facets of Web Services*. 19 December 2002, Gartner Group: Stamford.
60. Baltazar, B.H., *Web Services Impact Directory Standards Play Key Role*. eWeek, September 16, 2002.
61. B. Lheureux, F.K., *EDI to Web Services Networks: Contrasts in Convergence Transaction delivery networks and Web services networks may look different from traditional electronic data interchange value-added networks; however, their overall feature set is deceptively similar and consolidating rapidly*. 24 July 2002, Gartner: Stamford CT.
62. David Greenfield *Emerging Technology: XML - Racing Ahead?* Network Magazine, Jan 6, 2003.
63. Hiawatha Bray, G.S., *Akamai, IBM to launch service Teaming up to better deliver Web applications*. Boston Globe, 5/1/2003(5/1/2003).
64. Plummer, D. and R. Note. 24 January 2002, Gartner Group: Stamford Ct.
65. *"Downturn Hits IT Services Hard."* Information Week:, Mar 31, 2003.
66. Editor, *Six Keys to Building New Markets by Unleashing Disruptive Innovation*. Harvard Business School Week, Mar 10, 2003.
67. McGee, M.K., *Spreading the Word*. Information Week, May 27 2002. **May 27 2002**: p. 37-38.
68. Andrews, W., *CIO Update: Web Services Bring the Real-Time Enterprise Closer Web services can be the catalyst for the creation of service-oriented architectures. By letting*

applications access data across different devices, operating systems and locations, Web services. 19 February 2003, ICG.

69. Smith, D., *i-Deal Employs Web Services to Connect Broker/Dealers With the proper security measures, leading-edge Web services technology can link broker/dealers by enabling a platform-based approach to a financial application.* 9 July 2002, GartnerGroup: Stamford CT.

70. Gallagher, S., *Gotcha! Building An Application "Backbone".* April 1, 2003, Ziff Davis.

71. Jason Bloomberg, S.A., ZapThink,, *Ten Emerging Best Practices For Building SOAs.* ZapFlash
ZapThink's XML and Web Services Focused e-Newsletter, 02252003.

72. Schulte, R., *Predicts 2003: SOA Is Changing Software Service-oriented architectures and Web services are becoming mainstream. This will affect application design, many types of software products and even the nature of business-to-business value-added networks.* 9 December 2002, Gartner Group: Stamford Ct.

73. A. Mac Neela, B.W., *Storebrand Uses Web Services to Supply Financial Services*
Financial services provider Storebrand has expediently used Web services technology to track, verify and update pension plan records. Success in such an effort requires the alignment of enterprise business and IT strategies. 6 January 2003, Gartner Group: Stamford CT.

74. John Mahoney, M.R., *CIO 'Must Do' Resolutions for 2003;* 24 December 2002, Gartner Group: Stamford CT.


75. Vecchio, D., *MAINFRAME INTEGRATION*
Web Services and the Mainframe: Reality or Ridiculous?
Web services represent a new paradigm for connecting business processes independent of their underlying implementation. This approach can also be used to extend the life of legacy systems. 21 October 2002, Gartner Group: Stamford CT.

76. Porter, M.E., *On competition.* 1998, Boston: Harvard Business School Pub. vi, 485.

77. Hamel, G. and C.K. Prahalad, *Competing for the future.* 1994, Boston, Mass.: Harvard Business School Press. xv, 327.

78. Meyer, M., *The Power of Product Platforms: Building Value and Cost Leadership*. 1997: Free Press.
79. Gawer, A. and M.A. Cusumano, *Platform leadership : how Intel, Microsoft, and Cisco drive industry innovation*. 2002, Boston: Harvard Business School Press. xiv, 305.
80. Truog, D., *How The X Internet Will Communicate*. DECEMBER 2001, Forrester Reports: Cambridge MA.
81. Hagel, J., *Out of the box : strategies for achieving profits today and growth tomorrow through Web services*. 2002, Boston, Mass.: Harvard Business School Press. xix, 217.
82. Donald Lessard, R.M., *UNDERSTANDING AND MANAGING RISKS IN LARGE ENGINEERING PROJECTS*, in *Sloan Working Paper*. October 2001.
83. Reed, D.P., *That Sneaky Exponential-Beyond Metcalfe's Law to the Power of Community Building*. Context Magazine, Spring 1999.
84. *Business and Money*. Boston Globe, April 11, 2003.
85. Ricadela, A., *Road Map Redefined, Microsoft's new blueprint aims more squarely at business processes and collaboration*. Information Week, Oct. 14, 2002(Oct 14,2002): p. 70-74.

APPENDIX 1 MODEL



Market Scenario	Sensitivity	NPV New Business (\$ M)	NPV Total Business (\$ M)
Base Case	Base	99	244
Base Case	Rapid Commoditization	74	185
Base Case	Cross-sell high	407	667
Base Case	Cross-sell low	-15	101
Base Case	Niche market strategy	89	235
Base Case	Acquisition	248	417
Optimistic	Base	175	339
Optimistic	Rapid Commoditization effects	153	283
Optimistic	Cross-sell effectiveness high	501	768
Optimistic	Cross-sell effectiveness low	40	175
Web Services	Base	149	276
Web Services	Cross-sell effectiveness high	432	666
Web Services	Niche market strategy	143	271

APPENDIX 2. EIA

WORKSHEET

STARTUP: APPLICATIONS				
ITEM	YOUR TOTAL	DESCRIPTION	QUANTITY	EXAMPLE TOTAL
FOUNDATION SOFTWARE	\$ -			\$ 108,452
Application server	\$ -	IBM WebSphere Application Server Enterprise v.5.0	2	\$ 62,352
Messaging transport	\$ -	IBM WebSphere MQ	2	\$ 8,100
Database software	\$ -	Oracle 8i	1	\$ 38,000
DEVELOPMENT SOFTWARE	\$ -			\$ 32,398
Messaging	\$ -	IBM WebSphere Studio Enterprise Developer 5.0	4	\$ 27,000
Application runtime	\$ -	IBM WebSphere Development Runtime edition	2	\$ 5,398
Java	\$ -	Sun J2EE Developer Kit	4	\$ -
MESSAGING AND INTEGRATION	\$ -			\$ 247,050
Transaction management and integration	\$ -	IBM WebSphere Business Integration Server	2	\$ 222,750
Adapters for packaged applications	\$ -	IBM WebSphere Adapters	4	\$ 24,300
ENTERPRISE PORTAL	\$ -	WebSphere Portal Enable	2	\$ 129,600
Labor	\$ -			\$ 3,431,514
Business & Management	\$ -			\$ 298,542
Technology	\$ -			\$ 1,715,757
Consulting	\$ -			\$ 1,417,215
STARTUP COSTS	#REF!			\$ 3,949,014



OPERATIONS: APPLICATIONS				
ITEM	YOUR TOTAL	DESCRIPTION	QUANTITY	EXAMPLE TOTAL
SUPPORT FEES	\$ -			\$ 8,800
Database	\$ -			\$ 8,800
ANNUAL LICENSES	\$ -			\$ 632,701
Foundation	\$ -			\$ 108,452
Development environments	\$ -			\$ 17,999
Messaging/workflow/integration	\$ -			\$ 247,050
Portal	\$ -			\$ 259,200
Labor	\$ -			\$ 197,100
Business & Management	\$ -			\$ 32,850
Technology	\$ -			\$ 41,063
Consulting	\$ -			\$ 123,188
OPERATING COSTS (first year)	#REF!			\$ 838,601



	YOURS			EXAMPLE
TOTAL APPLICATIONS COSTS =	#REF!			\$ 4,787,615