The New Telecom Ecosystem:

Preparing for the New Ecosystem

Wedge Greene

LTC International

Date: 0705/2006

Article originally appeared as blog entries in InsideOut
http://www.ltcinternational.com/inside-out/

About LTC International Inc.

LTC International provides leading companies in the telecommunications and IT sectors with a unique level of service based on true subject matter expertise. Our Business Operations Architects® each have at least ten years of hands-on experience in service provider and IT intensive companies. Our consulting team has experience in all areas of business profit optimization, wireless and wireline communications, Internet services, as well as software and hardware planning, implementation and operations.

LTC has incorporated more than 1,000 years of first hand operating company and software application experience into our Business Management Toolkit. This comprehensive set of tools, guidelines, checklists, templates and training programs is designed to remove uncertainty and accelerate success for our clients.

http://www.ltcinternational.com

Expect Results®

THE NEW TELECOM ECOSYSTEM:

#1 – PREPARING FOR THE NEW ECOSYSTEM

Summary:

Collected from blogs written in the summer of 2006:

I discuss the operational challenges and competitive pressures facing traditional Telcos and new generation Service Providers as they prepare for the new telecom ecosystem. I look at operational efficiencies and inventive products as turn-around opportunities for Telcom. As a example, I provide some personal background history of our team's role in the creation of NGOSS (New Generation Operating Systems & Software) and how I came to be called the "Father of NGOSS". This article lays down the business drivers and environment that drives further discussions on Autonomic Communications in the companion article: Complex Systems and Autonomic Networks.

LTC INTERNATIONAL ®

Business Operations Architects ®

PREPARING FOR THE NEW TELECOM ECOSYSTEM

Part 1: Introduction

Original a series of blog articles by Wedge Greene.

> There is the short way that is the long way;
> There is the long way that is the short way.
> - Anonymous

I am sure that your experience is the same as mine: in the last five years the amount and percentage of family income devoted to paying for communication services has dramatically risen. Most of this increase is from our family's procurement of new services offered by communications companies. Costs for traditional POTS (Plain Old Telephone Service) are dramatically reduced, but this really amounts to a cost shift, for us from old to new service variants. These product shifts are exactly in line with the predictions made a decade ago by those tasked with setting the guiding strategy for telecoms.

Old Services	Replacement Service	New Services
o A local service surcharge for living in the country outside of the metro dialing area, o Metered long distance, o Second phone line [for the now discarded data modem]	o Un-metered AIN domestic service, o Unlimited domestic long distance, o Various pass along fees and taxes	o Mobile phone service for me and my spouse, o Mobile SMS data service, o Small-broadband home DSL service, o Various data subscription charges

$ Ca-ching €

If we include the current charges our family pays for non-telco services that may get bundled into future triple/quad play, IMS, and big-broadband to the home telcom services [cost of digital satellite, movie rentals, cost of home security monitoring, etc.], this amount is a healthy chunk of the family budget. And there are unrealized, new concept services I will gladly pay for when they arrive. So:

- *Why is the telecom industry in such bad shape?*
- *Why did it not live up to the lofty valuation expectations we bet our retirement funds on?*
- *What technical plans, business strategies and new service strategies will turn telco around?*

These are some of the fundamental questions we expect to frame in this series of articles discussing The New Telecom Ecosystem.

Series Roadmap

In early parts, I introduce the current telecom environment with its legion of problems. Then, as introduction, I provide some personal background history of our team's role in the creation of NGOSS (New Generation Operating Systems & Software).

In the articles to come over the next few months, we will discuss the operational challenges and competitive pressures facing traditional Telcos and new generation Service Providers as they prepare for the new telecom ecosystem. We'll look at operational efficiencies and inventive products as turn-around opportunities for Telcom.

We'll discuss the important issue of leadership, because visionary leadership seems to be lacking in telecom, or at least not visibly at the helm. Just recently, in the SBC takeover of AT&T, the technical leadership of AT&T's visionary SOA strategy did not survive the change. There is no surviving telecom executive practicing today that is a household name (and who was not a crook). But other successful industries all have these industry movers and shakers (Gates, Iacocca, Branson, Buffett, Ellison).

For some of us, it is not clear that telcom, as we now know it, will survive another decade. We believe there is no readily apparent sweet-spot to insure success. But we hope these articles can frame the story that new industry leadership must realize and sell. Somewhere out there is the executive that will rise to become a household name and the team that will be their tower.

Part 2: Constant Challenge – Constant Change

Getting hard data for the changes in the ratio of business revenue to telecom expenses is challenging. But it is clear that strategically, *telecommunication is more important to business, and to business survival, than ever before*. It is also clear that continuity and access to communications services are critical to the government in its foremost role of protecting its citizenship.

What the above change in my family product mix makes absolutely clear is that telecom is an industry undergoing accelerated change. Telecom is hitched to technology advancement. Technology is advancing rapidly and competitive forces drive everyone to participate in new advances. The most efficient company gets the big valuation. The rapidly growing company gets the biggest valuation. Reduction in customer subscribers leads to trashed valuations. So competitive forces also drive new service products. Customer churn is linked to competitive service advertisements and to ever cheaper pricing. Everyone is participating in an ongoing feedback loop of increased technology change and increased competition. *Where are the insightful strategies, careful procedures, and measured plans telecom was once famous in designing?* Telecom as an industry has moved from conservative monopoly, to ordered competition, to the free-for-all of the old west.

It is also clear that consumers and business will pay, and pay well, for new customer focused services. But this revenue is not assured; it must be earned. A decade ago, service quality and reliability was the overriding goal of nearly every telco; it amounted to a credo – it permeated the ranks of the corporations. Now, employees have learned to watch out for themselves before making sacrifices for the corporate vision. Strong customer brand loyalty has fallen to a feeling that telecoms cannot be completely trusted, certainly not trusted to put customers first. This is at least in part due to the now general awareness that Wall Street is the final judge of success – how good a job a telco is doing is 'how good is the stock price' and overall, the industry index. By that judge, telco is not doing well. Never mind how many of us personally lost from the telecom bubble – for the last 12 months, in a period when the market did well, Schwab puts *integrated telecommunication services* in the lowest strata of price-performance ranking.

What has not changed is that financial success still depends on good management, efficient operation, and customer-focused introduction of inventive, helpful products. So deep down I, and I expect this of most

telecom executives also, ask: Is it my fault? Am I contributing to the problem or the solution? I find that my colleagues passionately want to do right and to fix the mess telecom is in. *It just is not clear which is the winning strategy?* How do you navigate though the many voices saying we will make it right – just pick us?

Strategies	o Be more efficient o Faster introduction of new products o Higher quality and reliability o Cannibalization of media & entertainment services o Market specialization o Technical/Product specialization

Against this, many of the old-style technology targeted strategies are clearly past their expiration date. For example: "consolidated billing is what customers really want." I will argue that we are now past the point where "do more with less" can be a winning strategy. *We are no longer in the time of the "lean, mean operator"*. Today, the winning strategies will be grouped under: "do more with more." I do not mean just toss more bodies and dollars in existing approaches. Rather this is "do more with more smarts". However, a significant increase in investment and spending is needed for telecom to navigate these new waters. This must be a directed investment that is not about finding and eliminating the inefficient. It is about providing a common infrastructure and technical alignment that naturally eliminates duplication. It is about inventing autonomic processes that allow redirection of operational expenses from network focus toward direct customer interactions. It is about big investments now, to achieve order-of-magnitude reductions later - with leaps in productivity following.

A strategy of 'do more with less' coupled with the relentless competition enclosed in "lean and mean", with a lot of external market forces, have resulted in essentially a commodity market for services. In the worst cases this leads to treating customers as commodities. I believe, and hope to effectively argue in this series, that *the way out for telecom to again become a growth industry is to provide significant new services directed at customer needs, wants, and satisfaction.* And yes this might mean dressing up a chicken so that it seams better than every other chicken, as the Perdue marketing team achieved. But more fundamentally, it means being able to effectively deliver new products that people want, even when they do not yet know they want them.

Lastly, the telecom environment has fundamentally changed from what it was in the late 90's. It is no longer a balanced group of IXCs and RBOCs holding off edge incursions of CLECs and cable. Today it is a market showing specialization: territorial giants, wholesalers, business specialists, technical specialists (like mobile or wireless broadband), etc. Isenberg was right; the "stupid network" is achieved. Now we need to turn this around and build the "smart network." *And, as the Internet has shown us, it will be a cooperative network ecosystem.* With the proper common infrastructure for inter-working, this will not even be overt partnerships, but a free flow of resources and services, all resulting in real-time transactions inside the financial sector systems. So to recap – billing just goes away.

Part 3: A Short Personal History of NGOSS

I am not a "harbinger of doom." I did systematically summarize and then voice the multitude of problems facing OSS and BSS systems and Operations Support Centers at the turn of the century. However, *this was used as business justification for moving to new approaches* – specifically NGN networks. And I saw, with the conversion to new equipment, the opportunity to convert to a newly conceived telecom IT. At first this was specifically NGN-OSS (Next Generation Networks - Operations Support Systems): green-field OSS for the green-field networks that were coming. But all telecom services were in trouble and this soon became shortened to *Next Generation Operations Support Systems*, aka the first NGOSS acronym.

Nevertheless, the deep changes we advocated with NGOSS involved

- *fundamental shakeups in the vendors and the kinds of products they would offer,*
- *in the IT organizations of telecoms,*
- *in the weighting of investments in technology purchases.*

Many resisted this change, ultimately voicing their frustration in a cover article of New Economy magazine, "Worldcom's OSS Mystery Man". While I declined an interview at that time, those interests threatened by the vision we articulated spoke out. However, now I will speak up: offering my explanation of the sweeping changes I advocated. Simply, I saw a great number of critical problems and voiced these as corporate and industry issues. Agreeing with me, and seeing that technology alone would not repair telecom, the enlightened executives I worked for (the previously unaccredited pair Leo Cyr and Fred Briggs) *put everything on the table as resources for constructive change*. So we opportunistically explored for solutions in new IT technology, new switch technology, changes in operations procedures, changes in IT staffing and organization, and changes in vendor relationships.

Inventing the technical architecture and organization we needed to achieve was the easy part. It turned out that a systematic, out of the box, engineering exercise stepping back to take a completely fresh new look at the problems in total, explore emerging technological opportunities and apply system's science to look at the entire network-IT-user environment unveiled some neat solutions. Solving this was a team effort of a remarkable multi-disciplinary group that was assembled at MCI, soon augmented by WorldCom IT experts and some ace distributed-computing programmers from EDS. In assembling our solution we drew liberally from advanced computing teams in other industries, internet middleware and software designers (such as routing structures, DNS, RADIUS, Directories), and a 'somewhat rouge group' of experts from our competitors. These competitor architects were speakers at 'change oriented' conferences and included individuals from AT&T, Sprint, and several CLECS. Originally these conferences were side track sessions inside ATM and Internet conferences, but eventually around 1998, they got solo venues (via MarcusEvens and TeleStrategies conference companies). These conferences were the first to use the acronym NGOSS.

However, *charting a path from where the big, incumbent service provider telecoms were, to where they needed to go, was not easy*. Back in 1998, our MCI WorldCom advance architecture group approach was split into near term and long term phases: (1) start with what exists and integrate it in order to buy time for (2) inventing and building a new green-field approach. We also started parallel internal and external programs. Internally, start a seed NGN-OSS group that would create an integrated management system for new network products while designing a more efficient future telco operations system. Externally, seed our technical approaches in vendors who would create future products and provide architectural guidance via creating industry standards.

HOW TO ACHIEVE A NEW TELCO OPERATING SYSTEM, aka 1998, MCI WorldCom Advanced Technology team	Internal	External
Phase One: Near Term	Start with exiting technology and products and integrate these to buy time.	Provide architectural guidance via creating industry standards.
Phase Two: Far Term	Invent and build a new green-field approach.	Seed our technical approaches in vendors who would create future products.

In practice phase one involved using mostly modern, easily-integrated (rather than best-of-breed) vendor products in an integration architecture that knitted applications together. This approach was analogous to how phase one NGN was knitting existing networks together as a transition strategy while waiting for true NGN switches to finish development. For this NGN-OSS, we chose Publication/Subscription messaging technology (pub/sub) from the IETF as the framework glue. Phase two involved creating and developing a green-field, Service Oriented Architecture (SOA) that would provide a common base and approach for network service products, Business Support Systems (BSS), and OSS. We called this project NewWave. This internal program would re-engineer MCI WorldCom by creating a common technology platform and new processes. The external program would convince the vendor community to develop products for this NewWave style NGN-OSS. *In retrospect, our road map did not produce the vedor changes we anticipated.* Overall, our process for change was not successful; we were unable to achieve a full reengineering of OSS, BSS, and product systems into a common framework. Why?

In 1999, we chose the TeleManagement Forum (TMF) as the environment to advance our external program. The TMF provided an organized and systematic structure to broaden the support for these new approaches. The TMF had many characteristics and programs which led us to this conclusion. It had participated in most advances in OSS architecture over the decade it had existed and to this day continues to embody a culture of leading change. It was a venue where the entire telecom supply-side ecosystem participated: Service Providers, System Integrators, and hardware & software vendors. Plus it had three strong seed programs already in place: a common data definition program, the discussion of a component OSS program, but most notably, the Telecommunications' Operations Map (TOM). Additionally, several TMF board members, who had followed my talks on replacing TMN with NGOSS, saw value in this potential program. The response from proposing the TMF provide a home for NGOSS was dramatic and rapid. I joined their Board and the Board voted in the program in early 2000 and the membership began creating the teams and program in the Spring of 2000. As this was the start of the new century, the TMF Board changed "Next" to "New": the program name to *New Generation Operating Systems and Software*, the current name for the same acronym.

Since 2000 the TMF has led this activity with hundreds of extraordinary minds contributing to an evolving vision of NGOSS. While I articulated the original business drivers and contributed the 1st generation architecture, the program accelerated beyond my start with contributions from many companies. Today NGOSS is truly an industry product, while it resembles our original MCI architectural vision; it differs in many fundamental details, enough so it has become a different thing. Some of these differences are the result of implementation experience gained in the unique TMF Catalyst program. Some architectural differences come from the individual contributions of exceptional industry visionaries. Among them, with whom I worked closely, are David Raymer, Joel Fleck, and John Strassner. Including

the new eTOM provided a strong process orientation for NGOSS. Finally, the TMF staff provided a life-cycle vision of constant improvement and the leadership that made NGOSS an inclusive program.

But *differing business drivers account for the greatest departures the TMF NGOSS program has from our vision back in 1998*. The industry did not widely embrace changing the status quo to a new telecom ecosystem. Existing software vendors resisted the notion of OSS *components* – seeing this as requiring too much re-engineering of their existing products, but more importantly being a road to commoditization of OSS software. Of course, creating commodity OSS components with standard interfaces and functions was what the Service Providers needed. This feature was a fundamental strategy of the MCI WorldCom external change plan. Additionally, some equipment manufacturers resisted the development of inexpensive, autonomic management interfaces to their devices, seeing this as an incursion into their lucrative Element Manager product line. They implemented stalling strategies, slowing the release of NGOSS by generating strife in the working groups.

And there was no uniformity in the direction provided by Service Providers – as most were risk-adverse, they wished a slow and cautious approach to changing their existing base of OSS and BSS systems in which they had considerable investment. And Service Provider internal IT groups were organized around existing TMN FCAPS (Fault, Configuration, Accounting, Performance, & Security) functional domains - often complicated by vertical product lines. Silo OSS was duplicated with reinforcing silo IT organization. This cautious Service Provider wished to change only one thing at a time. This approach kept OSS components large and close to the existing definitions of management product groups: Element Managers, Network Managers, Service Managers, Manager of Manager, Inventory, Trouble Ticket Systems, and Billing Systems. Existing OSS software vendors supported and reinforced this view. These heritage product groups (which are not components) were integrated together via a Message Bus controlled by Work Flow products.

But message bus/work flow products never effectively solved the integration problem. While the architectural drawings and the messages themselves were much simpler, integration efforts were not really reduced that much over the old best-of-breed custom integrations. Those that still stick to this approach have replaced messages busses with web service connections. This approach seems functional. Verizon reports doing over a million web service transactions a day in their current OSS framework. But fundamental change was not achieved and consequently the strong results originally promised by a new OSS system have not yet been delivered.

For a brief time, a 19 company subgroup lead by my old MCI WorldCom team, and including a group of brand new distributed software vendors, an open-minded billing vendor, and some forward reaching equipment manufactures, attempted a solution that maintained a *service functional definition of Components*. This was called FineGrain NGOSS. I believe this group, in its Catalyst projects, provided the most remarkable ever example of integration, new service creation, and real-time autonomic management of modern network devices. An ecosystem of inter-working companies: infrastructure providers and a few key services reached commercial readiness. However, this approach did not get enough market acceptances to hold the key vendors together. I now believe the FineGrain NGOSS catalyst attempted too much in its demonstration of the capabilities of this approach, demonstrating twenty things instead of two or three – overwhelming spectators. It attempted to solve today's problems plus tomorrow's problems. Service Providers dismissed the problems which were not yet generating pain and the newly invented services being demonstrated. FineGrain NGOSS is no longer available and these concepts, once commercial, are now only a distant futuristic goal of the mainstream of NGOSS.

Instead an alternative approach based on creating standard Java interfaces between the heritage product classes gained more and more support. The OSS thru Java (OSS/J) group demonstrated a growing group of small but digestible advances. As more Telecom IT departments converted to Java Application Servers, the OSS/J group added new members. Now with the announcement by the TMF that OSS/J has moved out of "Java community" land and under TMF management, OSS/J has become the mainstream implementation method for NGOSS. The resulting program is being freshly branded as **Prospero** (It will get a full article treatment later). It currently has enough diverse backing from vendors, System Integrators, and Service Providers to succeed.

Still, the OSS/J architecture of directly-coupled Java interfaces between systems is no longer mainstream industry IT practice. Web services have demonstrated strong successes. OSS/J was moving in this direction. I expect we will soon see complementary web service versions of the original OSS/J interfaces as a part of the Prospero program. So, *NGOSS now is an implement-able standard with a vendor ecosystem available today.*

NGOSS was one of the original SOA architectural approaches. However, mainstream industry IT in groups like OASIS and the Global Grid Forum (GGF) have pushed the standard framework architecture of an SOA way past the heritage OSS, mid-NGOSS, and OMG component view. I hope current TMF NGOSS teams will monitor these organizations and push their advancements into new releases of the evolving NGOSS architecture. I will speak to this future in a later article, but now, back to where we started before this historical interlude.

* * *

Part 4: Heritage Problems

Do the original problems I identified in 1998 as business drivers for creating NGOSS still plague the telecom ecosystem? They do.

> "…[a survey responding telecom] CIO in a competitive service provider, pointed out that his list of challenges, the things that keep him awake at night, is pretty much the same as the list he would have drawn up ten years ago, when he worked for a large long-established telco in a non-competitive environment." - from "There's the Rub: Pain Points for Service Provider" By Barbara Lancaster in May 2006 Pipeline

What is different today is that these problems are now generally acknowledged. People know they exist and have worked at resolving them. However, the solution approaches taken are individually specific to each problem and therefore, piecemeal. Here is a list of these old problems:

- Bad data embedded in existing systems;
- Silo service development linked to product-specific OSS;
- After-the-fact OSS development, following technology implementations and product releases;
- Outdated, call based billing systems;
- Not-invented-here syndrome leading to telecom specific software products;
- Delays in service delivery (development and installation);
- Rapid network growth;
- New data communication technologies leading to new service products;
- Internet middleware & services deployed outside of OSS/BSS systems;

- Open competition & ILECS;
- MAD (Mergers and Acquisitions);
- Mobile communications technology;
- Technology churn in networks and software, so that nothing ever catches up and stabilizes.

In my next article I will enlarge on these Heritage problems speaking to what solutions have be advanced over the last five years. However, a fundamental tenant of my early business problem statement was that *these issues are synergistically linked*. Piecemeal approaches to solving them would not work. Did this bear out?

The Plethora of New Challenges

Complicating this mess of partially solved problems, there are new challenges [new technologies, new standards, new services, newly recognized problems, and external world changes] which are bearing down on telecom executives. We will address many of these in the coming months. Stay tuned and express your views via our Blog.

Part 5: Heritage Problems – Data

> "The total operating environment, for many service providers, is still not an enabler for success; it is a source of frustration."
>
> - Barbara Lancaster

I maintain that the problems we identified in 1998 as business drivers for creating NGOSS still plague the telecom ecosystem. Recapping the quote from our last article:

> "…[a survey responding telecom] CIO in a competitive service provider, pointed out that his list of challenges, the things that keep him awake at night, is pretty much the same as the list he would have drawn up ten years ago, when he worked for a large long-established telco in a non-competitive environment."
>
> - "There's the Rub: Pain Points for Service Provider" by Barbara Lancaster in May 2006 Pipeline

I am fairly sure if we sat down separately, this CIO and I would generate a different prioritized top 5 list of problems. But if we spoke together for 30 minutes, our lists would likely converge in agreement for the longer list of problems facing telecom. The older problems are now generally acknowledged but not solved. People know they exist and have worked incrementally at resolving them. Adding to them are problems new to our decade, and these tend to be whoppers:

- securing corporate data;
- network vulnerability to physical and cyber attacks;
- the staggering complexity of new services, networks, and the competitive environment;
- the demand for instant universal availability for any new service.

This new list itself is daunting. And it is the nature of things that new issues will be forthcoming tomorrow.

My specific message is that '*these problems are all synergistically linked*'. However strong the piecemeal solutions taken for each individual problem, the systematic web of interactions among these problems have not, and will not, allow piecemeal approaches, even in aggregate, to succeed in saving telecom companies.

In subsequent months I will expand on each of these new problem areas, as well as covering how we are doing on some of these older problems that I listed in the prior article. But now, let's review how we have approached one of the issues facing telecom in 1997: **DATA**.

Old Problem: Bad Data Embedded In Existing Systems.

Personal war story

Once upon a time as the millennium bug became a mandated top-down priority, a study at the service provider I worked for found that as many as 40% of all circuit references in the inventory system were not correct. A sizable part of these circuits did not even exist any more. But most upsetting, some of the non existent circuits were still being billed for every month. A balancing number of circuits that existed were not in the inventory system and had never been billed for. Some of this was due to 'correctable' errors, such as data entry mistakes in the manual data-entry systems. But some was the result of systematic process-driven errors as repair engineers in Network Operations Centers (NOCs), reacting within their 2 hour repair deadlines, switched customers from old circuits to new circuits. As the operations staff jumped to working the next failure in the alarm system, the old circuits became orphaned and the new circuits never got entered into inventory.

Data abnormalities like this were repeated in system after system, for every kind of data, in every operator. Competing solutions to correct all these bad data discoveries were advanced: (a) automate systems so one validated entry fed all the systems, (b) interconnect all the systems so changes in one part were fed to other parts, (c) make the network the "database of record" and resync all the systems by polling the network. But the theories struck reality: for (a) data formats were not consistent between systems, for (b) different system technologies and non-liner work flows made interconnection difficult and extremely expensive, for (c) the network did not contain the data necessary to link network elements to customers and orders – and the hardware vendors were not going to put that information burden on their switch/router products.

Personal perspective on current solutions to universal data exchange

To correct bad data spilling thru the systems, to integrate separate systems in common work flows, and for countless other technical drivers, not the least of which were regulated inter-carrier exchanges of information, we, as a collective of industry architects, realized '*we needed a common data language*.' Let me speak to the efforts that I was personally involved with (as an Executive, not a data architect). [Caution, technical acronyms ahead] At least three competing data standardization efforts were launched: the DMTF, the TeleManagement Forum (TMF) SIM, and the ITU-T tML.

Our team at WorldCom kicked off the tML project (telecommunications markup language, based on XML standard) and gave Ed White from our NGN-OSS team the job of establishing an XML vocabulary for inter-carrier exchanges. I meant this program to interface closely with other efforts like the TMF SIM, but the project was dominated by old TMN standards mavens who applied the older ITU CMIP data

model. This was a classic case of isolated design, a carry over for when national telecom labs generated all their own technology: "do it my way, the telecom way, because no one else understands our unique problems." So did they fix a failed standard, CMIP, by using a better description language, XML; or was this a case of perpetuating an outdated standard forward to plague a new generation. Our starting goal for tML, in order for it to become a specific technology implementation of the service side of the modeling efforts in the TMF, was never met. Admittedly, the TMF work was not ready in the time of need of the tML group and the CMIP model was there from the 80's. As the tML team progressed, the energetic and prolific team began looking further a field for data models to transcribe into tML. Eventually this team adopted information descriptions from a wide range of industries. And while it stands separate and apart from NGOSS; nevertheless, tML, now ITU-T standard M.3030, is worthwhile and greatly improved inter-carrier regulated data exchanges.

To overcome the different directions taken by the DMTF and the TMF SIM, I talked my friend John Strassner (creator of DEN) into leaving the DMTF and applying his considerable efforts and brilliance in data modeling to the TMF. Combining John and the existing strong TMF data team resulted in the creation of the TMF SID data model. With enough strong minds contributing, more companies got behind this effort. Today, t*he TMF SID is recognized as THE standard OSS data model.*

Good news: today there is one Telecom OSS data model standard: the TMF SID.

- **Challenge 1**: it is incomplete. Substantial parts are still locked up in the proprietary DEN information model and not yet part of the public TMF standard. Other parts are planned, but the contributions are not yet submitted.
- **Challenge 2**: [Caution: technical material ahead] by design, the implementation specifics of the bottom of every inheritance tree in the information model are left to the specific vendor implementer so that their specific attributes and extensions can be added. So just implementing the objects, makes every product instantiation different.
- **Challenge 3**: *in attempting to be both comprehensive and correct, the SID model is enormously complex*. Profiling the model into working data services will be a major effort. Strong implementation architectural control and inventive design is needed to profile an implementation of the model into software and data repositories.
- **Challenge 4**: the effort to build a common data service for telecom OSS has been superseded by the new world economics. Now a telecom ecosystem participant must inter-work with data (and processes, but that is another story) from a great number of external industries. The military has a different data model. The banking/finance industry has different models and data formats. The value cycle (logistics, shipping, supply-chain) has other formats. You get the point. *It will never be possible to have ONE model which all software can use to pass information around.*

But the need is still there. There is still bad data. There is still a great need to pass information and share it between systems. So do the System Integrators win and get tons of money doing countless custom integrations? No, *this is an opportunity for the shareholder/investor/customer focused executive to implement a synergistic ecosystem-based approach*. I believe a solution to all these data issues may exist in the development of de-facto data services. These are services like the web service mapping systems of Google and Microsoft. By opening up their mapping information via web service interfaces, these companies allowed 3rd parties to build applications using their mapping information. But we have a way to go before every industry and data domain has a *de facto* data-sharing supplier.

Part 6: "The database is evil"

I now believe a solution to service provider data issues (see prior posting) exists, not in a universal data model, but in the development of de-facto, data services exposed via web service interfaces. This new 'sharing approach' for data requires a different way of building systems. It is time to give up the monkey-on-our-back which is the database. Our systems architects and designers have an addiction to the database. Every implementation of any system must include a database. Then we relegate to the database all sorts of computation jobs – basically every association of information is done within the database and by the database. SQL is as much an association engine as any statistical product.

The *database was designed to associate information OVER TIME*. Putting information into a database stored the data so information put into a database at a later time could be linked to all the earlier information. The meal I had yesterday can be compared to the meal I have today. But the development of the relational structural model allows developers to organize structural data associations directly into how data was stored. Then SQL is used to retrieve organized information. [Ironically, today original the time-association job of databases is broken back out again as the Data Warehouse function.] It is bad programming design to use a database just to put data in, temporarily, it so that information can be pulled right back out in an organized form. Example: Customers are linked to orders and orders to parts inventory via a modeled relationship. Data is organized and stored according to this model. Then a direct data fetch via SQL and a business-transformation-process covert these relationships to a Purchase Order document. But that document is also broken back to parts and structural associations and thrown back into a database as a slew of records. Modern systems use data models and databases as general association engines. Object programming is designed to overcome the fractionalization of data into relational structural associations, but rarely is this realized by storing objects as native format data – too often the information is converted back into relational database format for persistence.

Yet databases are not actually efficient as procedural structural association engines. Following is a true parallel system test from around 2001. I architected both system concepts so I can trash the earlier design – but these original designs still dominate implementations.) The job of the two parallel systems in this test was to link usage data from the two endpoints of a frame relay circuit into a (circuit-call-based) billing record and consolidate several time measurements into a '24 hour bins in one day' usage record.

Database design:	Using standard design methods: distributed pollers gathered information from switches (who had only one end of the circuit and so only half the data) and put it into the database as the information was gathered. The frame relay virtual circuit endpoints were associated via the data model built into the database tables. And each circuit had many different time-points when records were gathered. So when the pollers were done a report pulled out the new records that included usage at both endpoints. This was batched for input to the billing system.
Service Design:	The competing new system was based on the model of a real-time aggregator. Pollers created events for each data instance and sent these to the aggregator receiver. The aggregator placed the information in logical bins based on knowledge of circuit ID and endpoints, structure which was feed as XML meta-data. As endpoints were associated and the time periods were complete, the billing records were published out in a steady stream - digestible by the billing system.

This problem was like a bunch of coins being sorted – both systems produced associations, but the aggregator had no database overhead and generated a steady stream of results. *In tests the aggregator was 1000x more efficient in processor/minutes.*

Part 7: Future solutions to data exchange between systems

". . . close cooperation across an ecosystem will stimulate new business designs, as companies redefine what they do and what they rely on others to do."
- IBM 2004 Global Information Outlook

For the many reasons I've shown, reaching a universal, integrated data model is so unlikely as not to be possible. In the future, to overcome the current data issues we have described, I maintain, *systems will be designed so that from the beginning, they do not own the data and information they process,* just use it for their need. Web services and other SOA implementation architectures, in their purist form, understand this. When all systems embody this, data will flow freely and data transforms will map all the data differences by understanding the data "sourcer" (kind-of-information, structural model, and reliability) and the current service needs. Notice I add reliability. No data in the future can be considered as absolutely true. All information is statistical in that it has a probability of being accurate. We could be getting accurate data or inaccurate data from the source. Perhaps one in a thousand times the information is inaccurate. Therefore, the data "sourcer" becomes important and the process of balancing data from many sources a normal activity.

This is the componentized data service architecture of the future. Data brokers will offer data via placing the meta-information into commonly accessible, well known Registries or alternately announce it via some advertising system (often word of mouth among designers who are members of common internet-based social network). In the best cases, common models (like SID) will be used to generate the meta-data schemas of these component data offerings. *Note that the data architect job drastically changes.*

Likely there will be charges for accessing this data. These micro-charges make it worth while for the data supplier to keep supplying data. The best data will come with a certification that it is reliable (along with the security certification). The consumer will use past knowledge and the assessments of reliability services (sort of like investment councilors) to determine how good the data is and how to treat the information. The systems manage this process via Policy, pre-loaded into the systems. Most of this public data will come from web services interfaces via web service protocols and be readable via XML schemas.

I expect some data schemas, will not be open web service interfaces, but will be private where companies will need to purchase or lease this data. Other data will come from closed-group interfaces, perhaps in secure, private protocols, in trust networks built up in authenticated, value-chain partner company groups. This information will be balanced by SLA agreements negotiated between the partners. (SLA designs will be a later topic)

Of course, some of these data sharing groups will just be the many departments inside a single company; others will be between sub-companies and master companies. Internally, departments can become profit centers that charge for the data they serve up. The more this data is used and the higher quality it is, the bigger their budget. Virtual profits can be distributed as local bonuses.

But much of the data will come from external corporate associations. In the future it will likely be cheaper, and provide greater risk reduction, for the service provider to go to an external vendor for most data. Evolving market dynamics could make most data free to use because the suppliers get their value from providing it via indifferent means, such as advertisements, good-will branding, etc. New system products will broker this data and provide transactions for incremental data use.

Part 8: Solving interlinked problems

Back to the synergistic interplay of the original problem list; my thesis – as regards the telecom issue of chronic bad data, *how are the complicating problem inter-relationships ameliorated in this new approach I outlined in Part 7?*

Problems:

- **OSS Silos & Rejection of Not Invented 'In My Department' Technology:**
  This ecological SOA solution to "bad" data does not attempt to own the data to keep it clean, or securely store the data to keep it reliable. It does not use data as a weapon to control the development and integration of systems. It does not need to, because it is involved in buying and selling information in a market. To succeed in this market, it needs to insure the data it supplies is as clean as possible and the supply is reliable.

- **MAD, Rapid Growth, Technology Churn**:
  The componentization of data and its accessibility, as linked to explicit and implicit SLA agreements, will allow client systems to be independent of who supplies the information. Mergers and divestitures will self correct with just data SLA renegotiations. New product creation just involves finding new suppliers for the needed information; design changes are easily assimilated via incremental system builds. As the data is externally sourced, and likely sourced

from multiple suppliers using countless numbers of data repositories, rapid and significant growth is accommodated by transactional requests across hundreds to thousands of segmented data stores.

- **Competition**:
  Sourcing and supplying data will become a central business arrangement for departments and companies. Besides beating a competitor via price, you can restrict that competitor from information. Or better, balance costs via selling the competitor what data they need; or setting up reciprocal data exchanges, based on common opportunity, not regulation.

- **Alien Internet Technology**:
  The architecture of this meta-ecology derives from basic Internet middleware. It is centrally incorporated into the design of the data exchange networks. (Yet, the system management problem must be solved another way.)

Taking this ecosystem direction may seem like a monumental task of design effort, require brilliant system re-engineering, and consume enormous implementation costs. I maintain that most of this will be evolutionary, incremental changes. Common industry models like the SID will be used to build the specific meta-schema for data exchange between specific subsets of the OSS systems environment. Conversion from the old database formats to these new structures will be largely isolated tasks undertaken as the data is opened-up to new consumers. The effort is incremental and can be assimilated in standard budgets. Those laborious, emotionally charged interdepartmental design meetings become interface exchange markets instead of political forums. *Entrepreneurs* will spring up to provide needed external data. Each new product introduced will save costs by going to this larger data market instead of sourcing data design and storage internally. Costs of new data builds will be partially offset by selling new information as well as new services. We already have internal telecom examples of this kind of ecosystem market: the schedule D, 800 number redirection; and another from the Internet: DNS URL name space allocation.

Data security is quite important in moving to an exchange market. Of course this market provides a potential nightmare for loss of control of personal information – way beyond Big Brother. Some common definition of public vs. private data may need to be legislated. But mostly market forces can control the details - if we generate a mechanism for controlling availability of private, personal data. For this I believe two types of institutions will spring up. First will be extensions of credit reporting companies that are regulated by government, yet also watched by their insurers, in order to reduce liability claims. Second will be a new breed of personal agent services that let users decide and input individual policy on how to make personal data available. These agents might also do active searches in the open networks to flag and harass uses of personal data which are not acceptable.

So I see the future not as a Wild West but as a *chaotic marketplace*. Here I use chaotic in a formal way. Engineering use of systems science and chaos mathematics will allow very stable market systems to arise. This will be the re-tasked job of the new Data Architect. The *line-architect* will be packaging data for offer in the open interfaces, using public meta-data definitions of structure; or those he knows of through his professional networks. The *data systems architect*, understanding chaotic systems, will build controls and agents which interface with the Internet or Extranet environments, using transactions and association groupings to push the ecology in the directions that are desired.

But what about the current data crisis?

In the meantime the extraordinarily wide dispersion of private data into so many different application data stores, each in its own company has laid open this information to data thieves. Breaches of the public trust by companies and agencies of the government have made the criminal accessibility of private information a weekly item in the evening news. Companies have been too lax at handling information and now they and their executives are paying for this, big time. It affects all of us; including me, including you. Last fall I received a letter from my old employer MCI, aka WorldCom, that an employee had let a laptop be stolen that contained my vital identity data and employment records. What was horrendous was that I had not worked there for over 3 years. Why was this data still in the possession of HR people on their laptops for their daily assignments? And it is not just companies making mistakes on information in their direct control. I also got a letter from Hotels.com that their accountant/system's integrator Ernst & Young had lost a laptop with data on my financial transactions with Hotels.com.

What response can corporations take to this calamity? Today! Because this problem is too immediate to wait for my future chaotic ecosystems. Well, MCI did the correct response. Soon after this data exposure, the company bought a site license to 'file encryption technology.' MCI mandated every user convert their desktop and laptop system to use the product. They managed the corporate-wide conversion to encrypted user resident files seemingly without major problems, in very little time. This class of product, file encryption managers, is available for purchase now and new and better products are springing up.

> "Ernst & Young, which has 30,000 laptops used by its highly mobile staff of consultants, is encrypting all contents on the computers, according to company spokesman Charlie Perkins."
> - Encryption a simple tool to protect data, Associated Press, By Stephen Manning

Every company should have this in use - so go do it NOW.

But this just protects the employee edge-system user-exposed data. What is being done about all that information stored in all those internal enterprise servers with their must-have application-specific databases, that all the many segmented IT groups generated as a part of their contribution to the mess that is telecom IT systems?

- Continued in "Complex Systems and Autonomic Networks"-