Can the Enterprise Keep Up with the All Digital Network?

Challenges and Strategies for Managing Data in the All Digital Enterprise

Bill Scheffler
Director, Broadband Services
SPATIAL*info* Inc.
9635 Maroon Circle, Suite 420
Englewood CO 80112 USA
+1 303 873-6880
Bill.Scheffler@spatialinfo.com
www.spatialinfo.com

Mark Zangari
Chief Technical Officer
SPATIAL*info* Inc.
9635 Maroon Circle, Suite 420
Englewood CO 80112 USA
+1 303 873-6880
Mark.Zangari@spatialinfo.com
www.spatialinfo.com

Introduction

It is now a common prediction that the all-digital, IP network is likely to foster a bourgeoning array of new services and network applications, and provide for the more cost effective delivery of current services. If this prediction turns out to be correct, MSOs will find themselves in a new competitive environment where the ability to quickly roll out newly developed, high-demand products will be essential to the health of the business.

This, of course, is not the first time the cable/broadband sector has found itself in this position. Indeed, most MSOs have either just completed, or are still in the process of, upgrading their once simple analog broadcast video services to more sophisticated digital services which now include, data, telephony, VOD, etc. This has forced most MSOs to:

- Upgrade their physical outside plant,
- Install new kinds of equipment in facilities such as headends and hubs,
- Modify their business and operational information systems to support the new network services.

To what extent can this recent experience be translated into the possibilities for innovation offered by an IP network?

This paper suggests that, unlike the recent past where the most time and resource consuming activity in offering new services was readying the network, in an IP network, less network re-engineering will be required, and it is most likely that developing the necessary operational and business support applications will become the bottleneck to large MSOs being able to commercially offer new services to their customers as quickly as possible.

Even the lessons from recent history point to this, with many of the BSS and OSS changes needed to support new services requiring months or often years to complete. However, since this has been consistent with the timeframes of the network upgrades, and since most MSOs have followed similar approaches, the "information system lag" has been masked somewhat. If IP lives up to its promise, however, there will be no mask behind which delays in supporting new services will be able to hide. In the IP future, successful MSOs will be those who are first to market with innovative new services, who offer higher quality customer service, and who have lower implementation costs. All of these will critically depend on the adaptability of the software systems which drive these functions.

This paper will address the key areas of information support, and the underlying data required for digital services, including:

- Geographical and logical addressing,
- Service inventory,
- Order processing and billing,
- Real time service provisioning and control, and
- Monitoring.

It is suggested that, rather than the current approach of building “bespoke” customized extensions to existing BSS and OSS platforms whenever the suite of service offerings changes, a far more efficient approach is to develop a flexible “meta model”, which defines all of the information needs of business operations, and can be quickly reconfigured as new services come on board. In the successful Digital Enterprise, months or years of software development associated with offering new kinds of products, or better ways of delivering products, must be compressed into weeks. In short, an approach is needed which does for software systems what IP does for the network itself; that is, provides a common platform which can easily manage new kinds of information. The latest software products are utilizing just such an approach; the challenge facing the Digital Enterprise, as discussed in this paper, is how to deploy them effectively.

Two key success factors stand out:

- High quality data, in terms of accuracy and completeness, and
- Tightly integrated, yet flexible systems.

Several benefits flow immediately upon achieving these:

- CSRs, having more and better information available, can more quickly and comprehensively attend to customer requests.
- Better information makes deployment of customer self-care (e.g. Web-based, IVR, etc.) more viable.
- Little or no overhead is wasted in transferring data between systems, or correlating data from different systems, as this becomes automated.
- Modifications to systems, such as those occurring when new services are introduced, can be made quickly and consistently from a single source.

In the successful Digital Enterprise, product innovation will not be hampered by “information system lag”, as the software platforms on which it relies will share similar advantages to those that have made IP so appealing.

Operating Cable Systems: Yesterday, Today and Tomorrow

It isn’t that many years ago since the Cable Industry delivered only Video entertainment services and only to residential customers.

The technologies, business processes, customer support staffs and operational support systems functioned in a very straight forward manner. Providing these services could be accomplished with Customer Service Representatives to take orders and schedule installations and service calls, Installers and Technicians to complete the installs and service calls, a billing system and addressable controller to bill, deliver and secure these services.

Today, Cable has greatly increased both the quantity and sophistication of Video entertainment services, and broadened its product offerings to include High Speed Internet/Data, Impulse Pay Per View, Video on Demand, and Telephony (both switched and Voice over Internet Protocol). Additionally, these services are offered to business/commercial customers as well as the traditional residential customers. Although the HFC network that carries these services from the Headends and Hubs to the customer premises has changed from the tree and branch coaxial network to a fiber-rich hub and spoke network, it remains a very straightforward means for distributing all of the offered services.

What have changed dramatically in the delivery of these new services are the technologies within headends and hubs, the customer premises, as well as business processes, customer support and fulfillment needs, and operational support systems.

Each of these services requires "smart" devices at the headends/hubs (Telephony switches, CMTS, HDT, HMS, and Video Servers) and at the customer premises (NIUs, Cable Modems, DCTs). All of these devices need to get information from and provide information to the operational support systems (billing, trouble ticketing, fault management, capacity management, etc.). In most present-day operations, the data transfers to, from and between these systems is cumbersome at best, and often includes steps that are partially (or sometimes totally) manual processes. Because of this, the time lag and labor expended in completing these process is greater than need be, and their ability to rapidly scale efficiently is questionable. But more importantly, the likelihood of providing incorrect service, billing incorrectly, or providing poor quality of service is significant.

Another operational change that exacerbates the delivery of these services is the movement of customer support staff away from being collocated with the delivery and operational support systems. In the past, the close physical proximity of support and operations personnel insured that the data needed to manage the provisioning, delivery and billing for services was accessible, even if it was housed in multiple systems and formats. Today the provisioning, fault management, capacity management, billing and customer care systems often need to be accessed by remote call centers, remote NOC technicians, etc. in addition to the onsite customer support staff. Often times all of the data that exists is not available to these off-site groups or not available in real time. This again increases the potential for providing incorrect service, billing incorrectly, or providing poor quality of service.

In addition, new attention is being paid to the extent to which marketing efforts can more effectively lead to higher revenues by overlaying traditional business intelligence and demographic data with information gathered from the network itself. This is particularly the case where limited capacity is available, and should be directed towards the highest revenue generating customers (e.g. commercial high speed data, residential VOD users, etc.)

Significant overhead is required to maintain and mine these siloed data repositories and to facilitate the exchange of pertinent data between them. Errors in billing or the delivery of services as well as any quality of service issues have a direct impact on the penetration levels of these services; if we can't get it right there are other providers for the customers to try. Furthermore, the ability to determine both the serviceability of a specific potential customer, and areas where an available service is being under-utilized is critical to maximize return on plant capital.

These issues have a direct impact on the profitability of the individual service offerings and the industry as a whole.

So, the contrast between the historical cable TV company and the modern MSO could hardly be more stark. As we have seen, today's MSO is a large, sophisticated provider of multiple parallel networked services. They operate in a dynamic, competitive environment over many demographics and geographies in the residential, commercial and community sectors. Yet, despite the complexity of their business, the operational culture of MSOs as an industry still bears many legacies dating back to the simpler days of community cable TV systems.

But it is changing. The recent widespread rollout of narrowcast digital services has resulted in many cable operators deploying more sophisticated OSS platforms. Most have also enhanced their billing system capabilities, not only in order to simply be able to charge for the new services, but also to provide them with the ability to bundle and package services in new and creative ways.

Looking to the future, the pace of change is only likely to increase as IP provides a more "abstract" network platform and speeds the implementation of new devices and new interactive services. To take full advantage of this, however, operators will need business and operational systems which are as adaptable as the digital network itself. Otherwise, while network engineers may be less bogged down preparing the network for new services than they used to, the "back office" will still face the same issues of making business and operational systems communicate with and control the devices used to deliver, monitor, and bill for service.

To fully realize the speed of deployment advantages the digital network offers, the ease of integrating the support systems—both among themselves, and with the network itself—must become comparable to plugging a pair of devices into an IP network. It is perhaps not surprising then that the software strategies for achieving this share many similarities with approach that has made IP successful.

Integration Strategies to Streamline Business Operations

The key difficulties with integrating the various systems which typically make up a cable MSO back-office/network operations environment are:

- Proprietary data structures and repositories.
- Closed architectures without open interfaces, or with proprietary interfaces, making it difficult for different applications to share data.
- "Monolithic" application architectures which cannot be easily distributed or scaled.
- Complicated (or non-existent) tools for extending, configuring, or customizing the system, making it difficult and expensive to adapt them to support new lines of business.

These points have directly led to the implementation and modification of these systems being costly and time consuming ventures, and their functionality has been difficult to extend in creative ways (for example, use live network data to produce a marketing map showing serviceable addresses color-coded by service take-up and degree of utilization, overlaid with demographics). But, until now, these drawbacks have largely been masked by the much, much greater cost and time that is expended in constructing the network itself. In the future however, where large-scale network re-engineering will not be required, even when radically new services are created, the ability to compete effectively with other modes of delivery will depend on an operator's ability to minimize time-to-market with new services. And this, in turn, will depend on their ability to rapidly deploy the support system infrastructure needed to provide the service, bill for it, and assure quality.

Fortunately, the situation just described is by no means unique to the cable industry, and much of the recent work in software engineering has been dedicated to solving exactly the issues just raised.

Models and Metamodels for Data and Processes

Modern software engineering developments read like alphabet soup: MDA, UML, CWM, MOF, XML, XMI, SOAP, WSDL... and the list goes on (and on). Yet, all these acronyms (some of which we shall define and explain below) are ultimately about defining a set of standards which will allow any system to easily interface with any other system, so long as both have followed the rules laid down by the same acronym.

The acronyms above define basically two kinds of rules:

- Modeling rules
- Interface rules,

and the two are closely related.

The first of the acronyms we listed in the "alphabet soup", MDA, stands for Model Driven Architecture, and is becoming a central tenet of current software-best practices. The concept behind it addresses the fact that, if different systems are to share data, then they must represent the objects that they share in a compatible fashion. To this end, one of the tasks that has traditionally consumed considerable time during the integration of different systems has been the definition of a common "vocabulary" of the objects each system represents. In addition, to share data effectively, each of the systems must agree on what properties or attributes are necessary to uniquely identify and describe a particular object.

A typical case is the sharing of address records, which are often separately maintained by engineering, billing, operations, marketing, etc. To the billing system, an address is normally uniquely identified by a house key; to engineering, addresses have a geographical location and an association to physical plant; in network operations, no physical address records may exist at all, and for marketing, the latitude and longitude of an address may be important for overlaying geographical demographic data sets. Clearly, each of these systems describes addresses in a completely different way, and there is no easy way that an address record in one system can be correlated to an address record in any other. Yet, just such a combination of data from different systems is essential to automate a process such as linking a customer service requests to the customer billing account, the plant required to fulfill the request, and the physical location of the customer in the event that an on-site visit is required. As the maturity of the digital network increases, then data integration of this kind will allow at least some service requests to be fulfilled using customer self-care (e.g. via a web site).

This level of integration can be most efficiently achieved if it is approached as an industry-wide, open standards initiative involving equipment vendors, network operators, software vendors and other stakeholders. Standards bodies both in the software industry (OMG, the Object Management Group is a prominent example) and also in certain industry sectors (such as utilities, as described below), have already demonstrated that it is possible to establish standardized models which will ensure that different systems will be able to share data. If the rules these lay down are adhered to, then the problem of defining a common vocabulary which is both spoken and understood by all the systems vanishes, and with it, the work and time currently needed to establish and maintain data model compatibility.

To return to our earlier example of address records, any system which needs to access or create information about an address would already implement the agreed framework for doing so, without any required customization or extension. This allows vendors of such systems, and the equipment which communicates with them, to create standard, "productized" interfaces for near "plug-and-play" data interconnectivity. (In fact, in the case of addresses, such standardization has partially occurred already, albeit in a limited way, in order to support E911—precisely because a large number of disparate systems must cooperate to deliver the service. However, the E911 address model is generally localized to supporting that service, and neither has the features suitable for modeling addresses for the business as a whole, nor is it generally deployed this way. However, the modeling approach suggested here can be thought of as extending the concept of the E911 addressing standard to all forms of information required to run a cable MSO).

Of equal importance to standardized data models are commonly agreed-to process models. Just as a standard data vocabulary is needed, systems which initiate or respond to events need to agree on what these events are, when they occur, what information they carry, etc. While the specific way event notification occurs, along with how information is passed between systems, are matters to be decided when defining interface protocols, defining the events themselves, the systems they affect, the information each carries, and their sequencing are all important aspects of defining a common business model.

Generally, in today's typical MSO environment, where event-driven integration occurs at all (flow-through provisioning is one common example), it is often achieved by creating custom interfaces for events which are highly specific to the systems being integrated. Developing these "from scratch" consumes time and resources, and repeating the process "from scratch" each time it is required for each operator and by each vendor only multiplies this further. By establishing a standardized process model with defined events, conditions, states, etc., the dynamic side of systems integration can reap similar efficiencies to those obtained by standardizing the data model.

Before closing this section, it is worth giving air to a potential objection that might be leveled at all of this standardization. Namely, that standardizing restricts flexibility. If a standards body dictates the way data is to be stored and managed, then how can each business tailor its systems to meet its own unique needs? Two common approaches to alleviate this are:

- Define layers of standards, rather than a single monolithic standard
- Define metamodels, rather than models.

both of which are widespread.

Under the first approach, various levels of compatibility can be achieved by defining a certain set of requirements at each level, which are more general at the "bottom" levels and become more restrictive towards the "top". Systems can be compliant up to a certain point (saving the effort of custom implementation up to that point), but can diverge from there. This method provides a mechanism for trading off the cost-effectiveness of standard functionality, against the tactical use of customization where required.

Under the second approach, rather than prescribing what data will be stored and how it will be organized (i.e. a model), a standard can specify the kinds of data which may be stored, and the methods for other systems to determine what that data is (i.e. a metamodel). This provides the required flexibility, but does so at some cost, as the interfaces between systems become slightly more complicated. Solving this issue is, however, currently one of the most active areas of research and development in the software industry.

By way of example, the utilities sector, faced with many of the same challenges, has begun a number of industry-wide initiatives aimed at standardizing data models and interfaces. Two of these are the Common Information Model (CIM), sponsored by the Electric Power Research Institute (EPRI, www.epri.com) and the MultiSpeak initiative (www.multispeak.org), sponsored by the National Rural Electric Cooperative Association. Further information about these is available at the web sites listed.

It's all in the Data

As a closing thought, it is worth mentioning that no matter how integrated the underlying software architectures become, systems are only as useful as the accuracy of the data they contain. Today, data accuracy is one of the principal factors limiting the extent to which systems can be utilized to their fullest potential, let alone enhanced by integration.

On closer examination, however, one of the principal reasons for data inaccuracy is itself the lack of integration. More precisely, data inaccuracy normally results from multiple copies of data residing in separate systems which do not have the ability to update one another. In this situation, when a change occurs—for example an address is corrected—that change must be separately applied to each of the systems in which that address is stored. Ensuring that this occurs is a challenge beyond even the most conscientious staff. As a result, an address update which may occur in a plant as-built is rarely propagated beyond that as-built. So, not only is an inconsistency introduced between the as-builts and other systems, such as billing and customer care, but some of these systems are now wrong. Furthermore, since there is no audit trail in place, when faced with inconsistent data, there is no way of determining which system, if any, is right.

In an integrated system scenario, instead of several disconnected copies of the data, there is one master copy of each record in a designated Database of Record. All updates to the data occur in this database, and are automatically distributed to other users of the data as required. So, if an update occurs in an as-built, that change will be automatically applied to any other systems which require or utilize that data.

Another key point in this situation is that data consistency is ensured without imposing additional workload on anyone. The real unseen benefit of an integrated system is that, as people go about doing their normal jobs, they are in fact maintaining the corporate database.

What should we do?

With the pace of change likely to increase, MSOs will no longer have the luxury of the lead times they have required thus far for their business and operational systems to keep up with the services they offer. We suggest here that a major portion of that lead time can be eliminated by the cable industry defining a set of standard metamodels and interfaces allowing the systems on which it relies to adapt to operational changes much more quickly.

What needs to be put in place to establish these standardized data and process models?

The software industry has already established the required metamodeling standards. Key among these is UML, the Unified Modeling Language, which allows both data models and processes to be described in a platform-independent way. It uses an intuitive, visual approach to represent objects and processes which is easily understood both inside and outside the software engineering community. It therefore provides an excellent means for subject matter experts and software professionals to communicate in a precise way, as well as a means for definitively documenting the types of standards discussed in this paper in a way that is not tied to any specific implementation. Examples of UML diagrams for defining relationships among plant support structures, and a process model for provisioning are shown below.

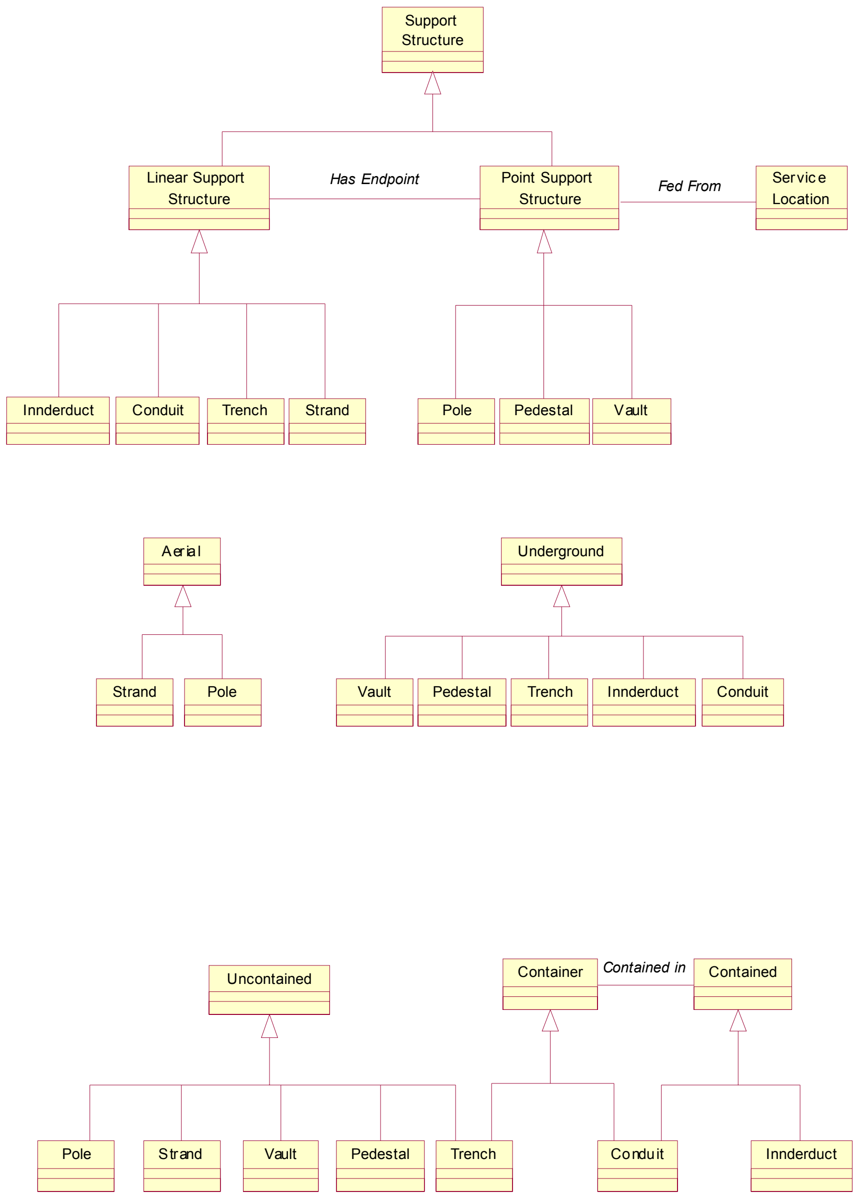

UML Class Diagram for example Support Structure model

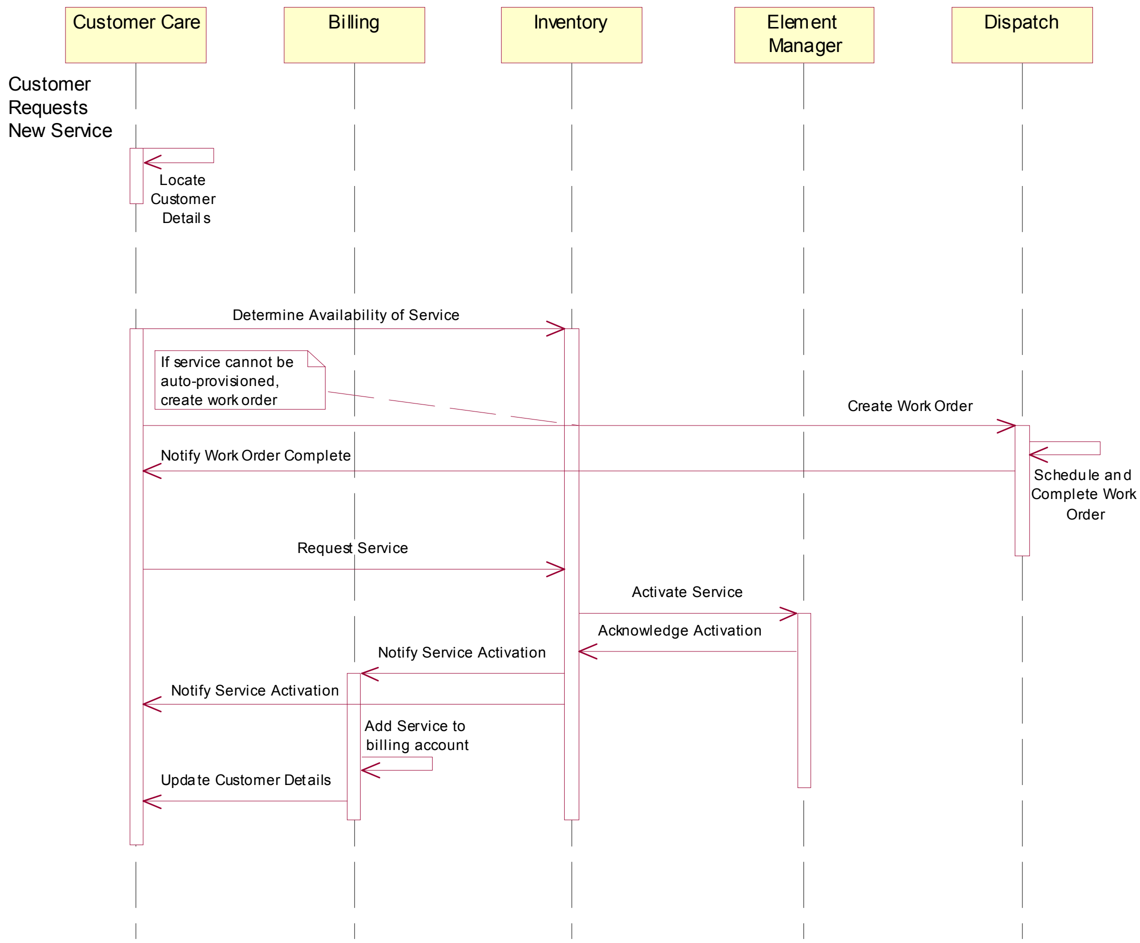

UML Sequence Diagram for example Provisioning Process

In addition to UML, a number of open interface protocols have also been developed, allowing systems to both exchange data in an open fashion, and to remotely execute services. Two of the emerging standards in this area are WSDL (Web Services Description Language) and SOAP (Simple Object Access Protocol) which use XML and are therefore independent of any platform and can use any kind of transport between systems (HTTP, files, SMTP e-mail, etc.).

To keep up with the digital networks they are offering, we suggest that cable operators, along with their equipment and software vendors, begin collaborating to define the software standards which will allow their businesses to be as adaptable as the new devices they are creating and deploying.