

Realtime Reconfiguration of Networks Through A Semantic Web

John C. Hoag, Ph.D.
 McClure School of Communication Systems Management
 Ohio University, Athens, OH 45701
 Hoagj@ohiou.edu

Abstract

Between all core and access data networks exists a broadband distribution layer, useful for interconnections of sites and perhaps coalitions of forces. Realtime deployment and reconfiguration of broadband wide-area transport is an important problem, and current solutions have modestly pursued admission controls or addition of functionality to routers – approaches that will limit interoperability and scalability. We are developing management tools to adapt network infrastructure in realtime by application-awareness. We propose to treat network operations and management transactions as language artifacts and solve the problem of service delivery symbolically and intelligently. In particular, our intent is traverse the following domain ontologies: from requirements to specifications, from specifications to configuration, and among various geographic coordinate systems. This has involved an ongoing search for telecommunications-related ontologies and the willingness to develop our own.

This research activity is aligned with the Worldwide Consortium for the Grid, with its interest in network centric operations, and the “Valued Information at the Right Time” initiative at the Naval Postgraduate School, whose published work involves adaptation of flight plans with respect to weather using semantic reasoning techniques.

This paper will discuss results and lessons learned from a software testbed for network management through semantic reasoning in largely an open source environment.

1. Introduction

The Internet’s TCP/IP protocol suite can be seen as a vertical stack of software modules, each layer wholly depending on services from the one below but not aware of its supporters’ internal

processes or status. This is the military’s *de facto* transportation infrastructure for all services (voice, data, video, etc.) and each layer contains the means for working around problems in the layers below rather diagnosis and repair of the root cause. The principle of independence between protocol layers has resulted in applications *not* aware of their networks and networks *not* aware of their applications.

This research is directed toward what has been avoided so far: automated remedies to solving problems at the level of infrastructure. Specifically, this research encompasses configuration (e.g, adding, modifying, deleting) of infrastructure elements in real time based on new collections of knowledge and new processing to support human network operators. Informed by Network Centric Enterprise Services goals and industry trends in general, this research has embraced the semantic web for organization and processing of information among the processing nodes that it requires.

Telecommunications is an ideal domain for demonstrating the semantic web since there are functionally equivalent approaches – SONET, MPLS, Ethernet, ISDN, and so on, or different satellites and transponders – that differ in name and description, assuming that they are available.

The single performance metric for this system is its business transaction cycle time for a service order, whose duration can range from hours to days in the current manual mode; a 45-day interval is accepted in industry for a circuit order when facilities are in place (i.e., no physical construction). As a consequence, this practical step of network design may take place only once in the life of a contract (e.g., 3-5 years) for stationary locations. The presence of new enterprise-facing business interfaces – implemented in XML in the best case – will permit this design-implement-operate cycle to be executed not only in minutes but repeatedly as an integral part of network management.

This project is adapting some emerging “open” tools to integrate an array of telecommunications network management and operations support systems, with particular interest in military use cases. A companion project is using the same logical approach to direct or select flight options in the presence of weather anomalies [1]. The historic motivation for this project was an investigation of the approach used to maintain stability in the U.S. electric transmission grid; in that case, topology is derived and flow models are balanced at hour intervals, with the system able to command both supply and demand [2]. Given that the grid involves many independent but equal firms, some of their system operators have implemented web services for wholesale power transactions.

This system’s activity currently lies outside the scope of NCES taxonomies, although telecommunications service and software providers are expending effort in this area. This paper only focuses on wide area network (WAN) services procured from commercial telecommunications providers, both terrestrial and satellite. The components of this system do not initially seem to take advantage of the spectrum of web service capabilities (e.g., discovery, service publishing, etc.).

II. Services

An implicit element of this realtime reconfiguration model is current knowledge of the applications supported by communications facilities. NCES identified the Global Information Grid to support all forms of web services in addition to common applications such as: [3]

- Session Management
- Text Messaging within a web conference
- Whiteboarding
- Application Sharing
- Audio (within a conference session)
- Presence and awareness (of others attending)
- File Transfer
- Voting and Polling
- Video, shared live or playback
- Recording
- Instant Messaging
- Chat rooms
- Virtual Spaces (for asynchronous collaboration)

Having established the identities of these applications, our next task is to identify highly-leveraged yet reliable means to efficiently acquire telecommunications resources. Domain ontology translations will enable three critical functions to take place in this model:

1. To translate the characteristics of applications in users’ parameters to engineering specifications.
2. To translate engineering specifications to sets of equivalent technologies to fulfill them.
3. To translate among alternate location references (address, latitude and longitude, industry coordinates, etc.)

Each application must be mapped to a service class, which may be a selection among few options for a particular telecommunications carrier: good, better, best; bronze, silver, gold; and so on. Each class is characterized by its latency, jitter, datarate burst tolerance, and error rate. Further, each application must be recognized by its geographic endpoints and datarate (i.e., bandwidth) in addition to its class.

For instance, Voice-over-IP would merit the highest class of service with least latency, jitter, and errors, with a datarate based on busy-hour peak traffic. With protocols to correct for transmission errors, Email can reside in the least class of service. Non-image sensor data could reside in an intermediate low-latency, low-bandwidth, high-priority class.

This is a substantial amount of information to be acquired and kept current, even with stationary network nodes. Moreover, while this information drives network design methodology, its derivation is not part of the task of a network operations center (NOC). Such traffic engineering is required for enterprise deployment of multiprotocol label switching (MPLS), the internationally-available infrastructure for “everything over IP.”

Consider the following maritime scenario, which applies to both commercial and military entities. At sea, reconfiguration of telecommunications services will take place when moving from the footprint of one satellite to another, affecting RF settings such as band and transponder; latency is constant, but transponder bandwidth will be reallocated among applications based on demand, such as when a videoconference takes place, or the other factors.

Navy SATCOM comprises several subsystems and numerous circuits over a variety of cryptographic, multiplexer, modem, and transceiver equipment. These subsystems arrived incrementally over time, and the complexity of the operations task is significant owing to disparate user interfaces. See [4] for a more complete description of these subsystems.

Depending on local regulations, functions carried over wireless transmission must be transitioned to terrestrial wireline providers. Installation of dedicated broadband facilities and reconfiguration of network equipment requires manual intervention by all parties from the order phase through implementation.

III. Two Forms of Network Management

This project previously reported on the two forms of network management and the gap between them [5]. In the context where the underlying network infrastructure is acquired through commercial providers, terrestrial and satellite, this gap is very relevant. Both forms of systems claim the title of network management.

The first form of network management is realtime monitoring of events “in-band” at the IP network layer. Network topology is externally constructed in this form, and elements of interest are individually identified. Further, the specific dependencies among related devices must be established and general “tuning” performed to *de-bounce* events that seem to naturally clear themselves fairly promptly.

The Simple Network Management Protocol (SNMP) provides the basis for site-specific information to be collected from product-specific Management Information Bases (MIBs), some of which are proprietary. The software network management system (NMS) may poll devices and accept *traps* of failure (and security) events. The NMS may be scripted to *walk* that device’s MIB tree to collect other parameters of interest. The ability to correlate events into a single alarm is the goal of site NMS developers. NMS software, which contains a local database of managed objects, may invoke the creation of a trouble ticket in a companion system which typically has its own representation of data.

Advanced installations may implement a hierarchy of NMS systems, and enterprises interested in monitoring service quality can invoke probes to monitor performance remotely.

In addition to the level of effort required to develop event correlations, the largest shortcoming of SNMP-based NMS systems is that they are only *open-loop* in nature. With effort, they can be refined to filter and correlate data, but they have little intrinsic diagnostic capability and no intrinsic ability to bring about change in the networks that they monitor.

Project team members have successfully been running Nagios and OpenNMS, both open source NMS packages that, respectively, use MRTG and RRDtool to facilitate data acquisition. OpenNMS is significant in that its internal data representations are in XML, and that through collaboration with another open project, it hands off fault management responsibility through a web service. [6]

Another form of network management at the business level comes by means of Operations Support Systems (OSS). Operated by telephone carriers, these systems provide data and transaction support for order entry, provisioning, field inventory, fault management, billing, among others. For the purposes of this project, inventory represents the set of all circuit elements at a component level, and provisioning represents the assignment of elements for a specific circuit. U.S. local exchange carriers develop their OSS systems internally (i.e., rather than COTS) to support “hundreds” of transaction types, only few of which must be “exposed” via web services or EDI.

The Telemangement Forum is a consortium of software vendors seeking to service established and emerging network providers. This consortium has developed a framework for a “next generation” of OSS, NGOSS, that defines business processes and data models for B2B interoperability [7]. NGOSS was developed with web services in mind, and it provides the means for vendors to certify compliance with the framework. Interestingly, the TMF has supported the development and demonstration of an open source project (OpenOSS). Its code is functional for fault management transactions, and it has been successfully integrated with OpenNMS.

IV: Bridging the Network Management Gap

The current intermediary between the two companion forms of network management systems is network operations staff, which must interpret the diagnosis of the NMS and find alternative infrastructure to fulfill the requirements of the applications at affected endpoints. This task complicated by the fact that customarily information about the users' requirements and the providers' inventories is not available in the NOC.

The foundation of our solution to this management problem is to invent a broad definition of network state and task a new node to maintain it.

- The set of application requirements by endpoint, aggregated for class of service
- The set of all network elements in provider(s) inventory
- The set of all provisioned resources (cross-reference of previous two items)
- The set of templates for configuring monitoring of all classes of elements

We have entitled this model, State-Based Network Management (SBNM).

We have entitled this model, State-Based Network Management (SBNM). At its core is a state engine that, among other tasks, seeks alternative network facilities from provider inventory when some failure. The state engine employs logic established by Hayes-Roth that exploits similarity in meaning across multiple ontologies in order to seek a solution. [8]

The following drawing identifies the participating agents in this model and represents the information flows between them. Nodes will be discussed below, and information flows are labeled by their protocol (e.g., XML or SQL).

In Figure 1, the central node is the SBNM state engine that employs VIRT reasoning. The right box is an agent that contains knowledge about requirements (applications, CoS, locations, etc.). The leftmost box represents some provider's OSS suite, specifically the inventory and service activation functions. Attached, above, to the state engine is a relational database that fulfills the semantic search for new infrastructure. Very significantly, the bottom box represents a self-contained NMS whose only external dialogs are to announce a confirmed failure and to accept configuration data for something new to monitor.

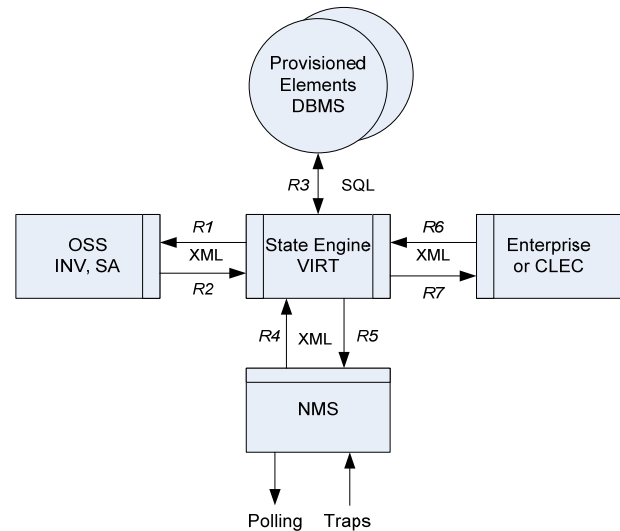


Fig. 1. Graphical representation of State-Based Network Management

Two cases exemplify typical and exception processing. In the typical case, a request for new connectivity for some application will appear; the user's agent will translate this into CoS and bandwidth requirements, which are sent (R6) to the processor. The processor will register the request in its database (3), and then semantically search the database for equivalent solution types to fulfill the requirements. For instance, a 512 Kbps symmetric circuit with, say, 25 ms one-way latency can be fulfilled by:

- A fractional T-1 circuit
- Two bonded ISDN BRIs
- A pair of frame relay or ATM virtual circuits with (CIR, Bc, Be) stated
- A point-to-point metro Ethernet
- An MPLS VLAN subject to CoS

These solutions can be ranked, then sought within provider inventory iteratively (R1 and R2). The best solution is ordered, provisioned, and activated, then posted to the database. The NMS is notified (R5) of some new facility to monitor, then the user (R7) is notified of service availability.

Alternatively, the process can be invoked by the NMS upon confirmation of an element failure (R4). Its application information is retrieved (R3) and the best available alternative solution is acquired (R1, R2) and implemented (R3, R5, and R7). We foresee that solutions are offered to the NOC staff for selection.

While this schematic represents one NOC, the concept is easily extended to a grid setting. Each SBNM node maintains its local state based on local requirements. In this regard, rather than using a commercial provider, SBNM nodes can cooperate to pass traffic among themselves. In this case, SBNM nodes could adopt the commercial provider's ontology and service oriented architecture.

V. Results and Assessment

This paper has defined a problem of practical significance to military interest and offered the architecture for a conservative and workable solution. Research progress is aided by the presence of open source efforts in all aspects of network management as well, through our use of Protégé, of ontology development. This paper has offered a sensible metric as well as a range of use cases, and anticipates a two-orders-of-magnitude improvement in service intervals when in service.

Ongoing work includes the integration of open source tools and participation in further architecture and validation efforts.

References

- [1] Hayes-Roth, Frederick, "Model-Based Communications Networks," Version 1.1, manuscript, 2004.
- [2] Hoag, John C., "Telecommunications in Complex Systems Failure: The Electricity Blackout of August 15, 2004," paper presented at the 12th International Conference on the Communication Systems Management, Monterey, CA 2004.
- [3] Defense Information Systems Agency, "Net-Centric Enterprise Services Collaboration Service Request for Quotation, Statement of Objectives," Attachment 2, December 16, 2005
- [4] Kim, John C., and E.I. Muehlendorf, Naval Shipboard Communication Systems, Prentice-Hall, 1995.
- [5] Hoag, John and Chris Gunderson, "State-Based Network Infrastructure Allocation," presented at IEEE MILCOM (Military Communications) Conference, October, 2005.
- [6] Gallen, Craig, and Jeffrey Reeve, "Investigating the Feasibility of Open Development of Operations Support Systems," in Proceedings of the 9th IFIP/IEEE International Conference on Integrated Network Management, 2005.
- [7] Reilly, John P., and M.J. Creaner, NGOSS Distilled, Cambridge, U.K.: The Lean Corporation, 2005.
- [8] Hayes-Roth, Richard, "Model-Based Communication Networks and VIRT," Presented to the 10th International Command and Control Research and Technology Symposium, 2005.

John C. Hoag is an Assistant Professor in the McClure School of Communication Systems Management at Ohio University, Athens, OH, where his research interests are network management and network performance. Dr. Hoag is an Institute Associate of the National Regulatory Research Institute. His degrees are in Computer Science and Industrial and Systems Engineering, with an M.S. and Ph.D. from The Ohio State University. Dr. Hoag is a Senior Member of the IEEE, belonging also the IEEE Communications and Systems, Man, and Cybernetics Societies.