PipelinePub.com

June, 2006 | Volume 3, Issue 2

Supercharging Inventory Management

By Costa Constantakis & Sergio Pellizzari Nakina Systems

Service providers in the telecom and data networking industry spend three times more to operate and secure their networks than they do on capital investments to buy networking equipment. It's little wonder that they are showing a keen interest in reducing operating expenditures (OPEX) through improved efficiencies and better use of available Network Management / OSS software tools.

When discussing such OSS tools, particularly in the real world of network resource management, most of us commonly refer to inventory management and related applications such as provisioning and service activation. But a strong case for network inventory management in particular has been made in service provider circles and is now widely accepted operational doctrine. Leading vendors like Telcordia, Cramer Systems, MetaSolv Software and NetCracker are filling the needs in this realm adeptly with feature-rich offerings in their respective inventory management applications.

Yet in spite of the strength of inventory management applications available on the market today, service providers are still consistently frustrated by inconsistencies in their databases of record that hamper their efforts to automate critical processes like service activation.

The question, then, remains: Why is that?

Garbage in, Garbage out

The inherent weakness in Inventory Management systems lies not in the applications themselves but in the way data is fed into them.

Service providers today largely use manual processes to populate and maintain database information that relates to the network itself. This data is intended to represent real-time information on the network infrastructure. In many cases, these databases exist as disparate "islands" of individually-owned data in spreadsheets or other rudimentary database formats. It is not unusual for a service provider's inventory to be tracked using thousands of individual spreadsheets residing on the computer desktops of geographically dispersed individuals, with no particular central coordination to synchronize or update the data in a central location. These data stores are manual, inconsistent, and largely inaccessible by individuals or "northbound" software applications for which such data would be tremendously valuable, such as inventory management systems.

The case for consolidating these databases onto a single, robust inventory management platform has been made very effectively by solution vendors in the industry.

The case for automating the data input process is equally strong. The process, commonly referred to as Auto-Discovery or Network Discovery, has been talked about extensively but

only recently has it begun being taken very seriously. Major players like Telcordia (through its acquisition of Granite) and Cramer Systems (through its recent acquisition of T-Soft) are publicly proclaiming their entry into this space.

Without Network Discovery, service providers are faced with an inability to accurately "see" all of the network elements deployed in their networks at any given time. Furthermore, the ability to track changes to each network element, along with information pertaining to its configuration, interconnections, and available facilities, has been an elusive goal in practice. With changes being made day in and day out in a service provider's network, keeping track of changes using the manual processes common today is nearly impossible. Yet without an accurate representation of these basic pieces of information, network administrators will continue to face an uphill battle in trying to automate any part of their service activation and provisioning or network engineering activities.

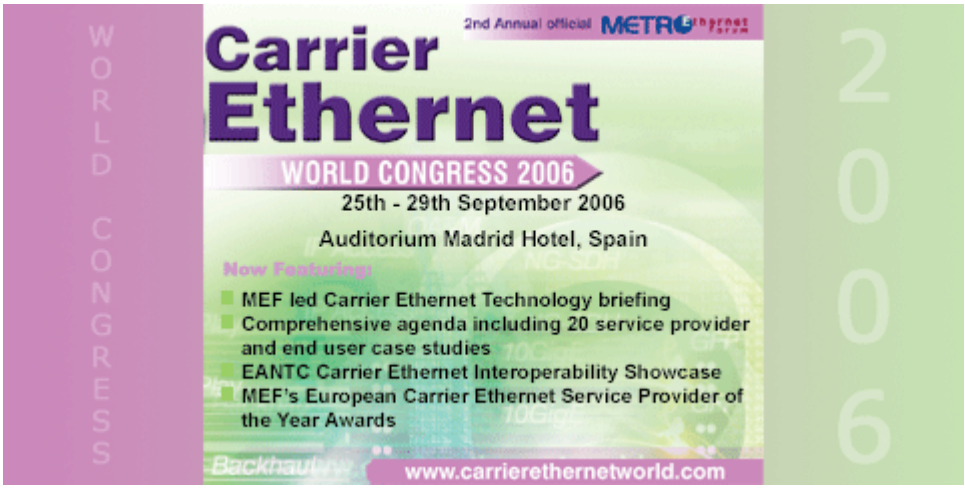

Scratching further beneath the surface, consider some related problems that arise from this approach.

$ **Stranded assets that result in unnecessary capital spending.** Networks are over-provisioned both by design and due to poor knowledge of which facilities are actually available in a network. In one study, industry research firm OVUM-RHK estimated that up to $10 billion in network assets in North America alone could be recovered and reused by deploying better software tools for database reconciliation in real-time. This leads to network equipment purchases that would otherwise be unnecessary. RHK further estimated that a typical network inventory database will inaccurately represent installed network resources by 20% to 40%. The problem is one that tends to be exacerbated over time as network operations personnel pressured for time find themselves too busy to record changes in the database. Unbeknownst to network operators, these assets are actually available for use, but the database misleads them into believing otherwise. Service providers are materially impacted, needless to say, as they find themselves buying networking equipment that would otherwise not be necessary, cutting substantially into corporate profits.

$ **Long provisioning cycles that negatively impact customer satisfaction.** Inaccurate data in the inventory management system, producing an inaccurate view of the state of the network, requires the provisioning staff to conduct time-consuming resource verification and database reconciliation with the network database before new services can be provisioned

and activated. These operations are time-consuming and labor-intensive, leading to costly lengthening of provisioning cycles and delays in realizing new revenue. More importantly, the provisioning delays often lead to customer annoyance and, by extension, the notorious churn rate in the industry. Service providers who show themselves capable of meeting customer needs expeditiously will enjoy a strong competitive advantage in the years to come.

$ **Trouble Resolution and Service Assurance.** Unmanaged faults derive from poor topological visibility. Without solid knowledge of the relationships between elements and facilities, it is nearly impossible to understand what alarms affect which elements and services. For example, a system alarm signaling the failure of a network element or group of elements has limited benefit if the network operators cannot pinpoint the elements that need to be addressed or the services that are affected by the network failure. In fact, an often-made request by service providers is to have the capability to visualize the fault status of a network element contextually, i.e., superimposed on the graphical topology display. This is preferable to a simple prioritized, filtered list of alarms and helps expedite troubleshooting.

$ **"Garbage in, Garbage Out" – hampering automation with bad data.** Northbound software applications, such as Inventory Management and Customer Care applications, provide network operators with actionable information that is only as good as the data fed into the software. The repercussions of using data that is substantively inaccurate can be dramatic. Salient examples include revenue leakage due to billing systems that do not reflect all customer services as well as terminated services not being removed from the network to free up valuable resources.

Clearly then, service providers need a more effective, manageable approach to tracking equipment and facilities in their networks if they seek to reduce their operating and capital expenditures while extending market share.

Network Discovery – Finding the Right Solution

Network discovery engines are gradually going mainstream, for all the reasons explained above. But what constitutes a good discovery engine? That is, what characterizes an engine that provides relevant and accurate data to an inventory database reconciliation process in a way that meets the operational needs of the service provider?

The following are key characteristics that differentiate true network discovery solutions from the pretenders.

1. True Discovery, not "Give-Me-Data-and-I'll-Confirm". Some discovery solutions offer good reconciliation capabilities but in truth require the network operator to provide the majority of IP or OSI addresses in the network before checking the accuracy of data provided and digging deeper for lower level attributes. A true network discovery engine takes in some minimal data (e.g, gateway node addresses) and discovers the physical and logical attributes of the network.
2.
3. Multi-technology, multi-vendor capability. A handful of network discovery engines work well within some areas of the network, such as IP networks, but fall short when it comes to the optical network layer, for example. Others only discover next generation networking equipment in optical networks, dismissing more than 75 percent of deployed legacy networks (especially in the optical layer, where older vintages don't provide data on relationships with neighboring nodes). A good discovery engine offers the capability to cover most of the deployed network that the service provider is looking to discover and track.

4. Scalability and Robustness. Enterprise-level solutions do not adequately meet the needs of Tier 1 carriers for scalability (with tens of thousands of network devices) and robustness. Most discovery engines built specifically for carrier networks handle this requirement quite well, with tiered architectures and server redundancies enabling load-balancing, etc., to maximize availability and prevent failures.

5. Leverage Existing Standards, but Look for Flexibility Foremost. Standards are important – but the reality is very few are universally adopted by all service providers. Therefore, a network operator needs to consider whether its chosen discovery engine supports the interfaces it has implemented. For example, if a TMF814 or MTOSI interface has been adopted, then the discovery engine selected needs to accommodate those particular interface requirements. In truth, a well-designed discovery engine will have a service-oriented architecture and should be easily customized to interface with the applications requiring the discovery data output (e.g., the inventory management reconciliation process).

6. Flexible Usage Parameters. Network operators need flexibility in configuring the usage parameters of the discovery engine. For example, how often a network discovery sweep should be activated, or how much of the available network resources the sweep should be allowed to consume (e.g., bandwidth).

Conclusion

In its latest market report on Inventory Management, Dittberner Associates, a leading OSS market research firm, wrote that it "… sees some opportunity for inventory vendors to partner with or acquire network reconciliation and discovery vendors. Network discovery is an important inventory verifying technology that makes the inventory solution all the stronger."

In the same study, Dittberner suggested the inevitability of network discovery's emergence by way of relating how expensive it is for telcos to physically audit their network assets: The study noted one example of the "… one telecom CIO [who] estimated that it cost him about $10 million to audit his network facilities, comprising 34 switches across 40 markets".

Network Discovery will be a true differentiator for inventory management solution providers and independent software vendors moving forward. Choosing an inventory management solution with a robust discovery engine will enable service providers to build competitive differentiation into the level of service they provide to their customers, through markedly improved turnaround times on requests for service activation and an improved ability to properly engineer their networks, based on accurate network data.

Discovery is the key. And it's happening now.