Alcatel·Lucent

STRATEGIC WHITE PAPER

Strategies for Supporting Next-Generation Networks

Proactive, predictive and preventive

As service providers compete to roll out advanced end-user services over next-generation networks (NGNs), they must consistently meet accelerating customer performance and service expectations to maintain their competitive edge. A key component in balancing these efforts is the development of a new maintenance model that goes beyond conventional product-level support.

This white paper dimensions the two-stage evolution from today's prevalent product-level support model to network-level support and ultimately to service-level support. It then details the four key ingredients of the service-level support methodology and reveals and how this methodology provides the tools for seamless end-to-end proactive monitoring and maintenance that optimize service availability, reduce operating costs and allow service providers to focus on service enhancements for revenue growth.

Table of contents

Introduction

These are exciting yet challenging times for service providers. They are rapidly building next-generation networks to deliver new end-user services in a highly competitive market. They are growing their service bundles with advanced applications such as peer-to-peer video, business/consumer voice over IP (VoIP), IPTV, gaming, and more.

The greatest challenge is delivering the services customers want with the quality of service they demand. To do so, service providers must proactively monitor and maintain their next generation networks to operate at peak performance. That's the key to delivering superior service, successfully rolling out new applications, and maintaining a high level of customer satisfaction.

As the traditional TDM circuit, data and mobile networks converge with a rich set of service interaction capabilities in the next generation networks, operations and maintenance strategies naturally evolve, to complement the conventional operations and product-level maintenance services. These new maintenance strategies include the network-level model and the service-level model.

Typically, the network-level model has the following objectives and benefits:

- Improve time to market for next generation network's end-user services by augmenting the service provider's staff with an experienced workforce thus reducing initial and subsequent training requirements.
- Reduce trouble resolution time by providing seamless support for network and product-level issues — from notification to trouble analysis & isolation to resolution.

The network-level model will quickly evolve to service-level maintenance, with the following characteristics:

- Predictive modeling helps to plan and deliver the desired service availability targets and allows service providers to take corrective actions before customers complain.
- Near real-time service monitoring against key quality indicators helps identify potential problems before they can cause real trouble.
- Proactive response empowers providers to eliminate potential trouble spots before they can become real threats to the services delivered.

CUSTOMER SATISFACTION: MORE IMPORTANT THAN EVER

As innovative services become widely available, consumer expectations about service quality (QoS) increase dramatically as well. A few years ago, a dropped mobile call or failed call attempt would have been tolerated while a new service was being rolled out. Today, the bar is raised as consumers learn to depend on each new service for business and personal needs. Consumers expect new services to work at least as well as, if not better than, those they currently use and pay for. Simply put, consumers want excellent service — anytime, anywhere, on any device.

Providers that don't pay close attention to customer QoS do so at their own peril. Consumers have already made it clear that if they're not happy with a service or a provider, they will switch to a competitor. That drives up churn and drives down billable revenue.

The market is crowded with competitors offering very similar service bundles. Price-points are all but identical. Brand loyalty can be fleeting. Service quality has become one of the key competitive differentiators. Maintaining service quality for millions of subscribers across the service lifecycle is a business imperative.

The evolving maintenance model

Service providers are all too aware of the competitive landscape. They would like to be able to deliver guaranteed service level agreements (SLAs) to their customers, with high service quality and availability. These guarantees must in turn be supported by the vendors who sell the new technological solutions to the service providers.

The problem: many service providers currently rely on product-level maintenance methodology. This approach focuses only on the health of individual network elements. Because it does not address an end-to-end view of application or service availability, the service providers have limited visibility of potential trouble areas, or how their services are performing until a frustrated customer calls to complain.

In addition, today's product-level support drives service providers to establish a separate point of contact with every vendor whose products they have deployed. As a result, end-to-end service-level problem resolution can be cumbersome and time consuming — and the longer a problem persists, the unhappier end-users can become.

Most important: A product-level approach is, by definition, limited to the individual products in the network. **But not all the problems on a next-generation network can be isolated to a specific product.** A connection failure or dropped session, for instance, isn't necessarily tied to a specific product. It could be caused by congestion or a session timeout. Soft failures cannot always be traced back to a particular product.

For these reasons a new maintenance model is needed: one that provides a single-point-of-contact for supporting the deployed end-to-end network solution, and delivers both reactive product-level and proactive service availability support. This new approach offers an end-to-end view of a provider's network to meet service level expectations and to be more responsive to customer needs.

This new maintenance model will evolve along two axes: 1) moving from reactive (break/fix) to proactive and, 2) from product-level support, through network-level support to service-level support.

Ultimately, the network-level model will evolve into a service-level solution that continually and proactively monitors a provider's application and service availability. It will make use of performance statistics, key performance indicators (KPIs), and key quality indicators (KQIs) to enhance both services and networks, and to optimize service performance and reliability end-to-end. It also will use predictive modeling to anticipate and avoid problems. Similarly, near-real-time proactive monitoring detects problems and allows support to correct them before they become service impacting. This service-level model will increase the stability of service provider networks, thus enabling them to roll out new services more rapidly and drive increased revenues.

Product-level support

Many service providers currently subscribe to maintenance technical support offers for all the products in their network across all the vendors. These support agreements typically include support-level KPIs, based on the acceptable respond, restore and resolve times, to ensure that the OEM vendor will react to a problem in a timely manner. However, the responsibility for identifying network and service level issues is with the service providers. It requires the service providers to maintain separate points of contact with each of their equipment vendors and manage any interoperability issues among those vendor's products to resolve end-to-end service problems.

With the conventional maintenance service model, as the service providers react to the problems in the network, they could easily spend over an hour attempting to isolate the root cause of a problem — *before they know which vendor to contact*. The process involves collecting and interpreting alarms and fault data, then correlating network performance data with configuration of the network and

application. In some cases it might be very difficult for a service provider to identify a network-level problem with only product-level visibility, which delays service restoration.

Once a service provider has the relevant information needed to identify where a potential problem has occurred, the Service Provider contacts the specific vendor for product related issues. That, in theory, should be the beginning of the end of that particular problem. However, sometimes there are disagreements among vendors as to where a problem has occurred. The identification of product root cause can be difficult when a service operates across several products. During this back-and-forth discussion between vendors, the service provider loses precious time as it attempts to get to the cause of the problem and correct it.

End-users, meanwhile, don't know or care anything about this process. What they want is for their interrupted service to be restored. Now!

THE BREAK/FIX PARADIGM

Today, most Service Providers leverage the conventional product level maintenance in a reactive paradigm. The provider and vendor focus on fixing what's broken, which is why this type of approach is typically known as a "break/fix" model. Basic product level maintenance support does not offer the tools to be able to address problems proactively, rather, it addresses correcting problems after they have occurred.

To be fair, the product level approach has worked successfully — especially, with the legacy circuit-switched network that involves fewer vendors with discrete functions. There are several reasons for this. First, legacy networks are typically not multi-vendor but limited to just a few vendors who supplied products to cover specific functionality. As a result, a service provider is dealing only with a few vendors and a few support agreements. For example, a circuit switch as a single network element provided both transport, call control and service application. If the switch performed well, then the service application did as well. It was clear where to turn to if a customer service failed — the circuit switch supporting the customer.

In contrast, the converged and distributed next-generation networks involve many more vendors, each of whom may have multiple products of different functions. IP-based services are inherently distributed, with different components providing parts of the application, control, and media layers. Now, a provider must trouble shoot a service problem across many types of products from dozens of vendors, because no single component by itself can determine whether the service is being successfully delivered to the end-subscriber.

Additionally, in the past, service providers built and refined their legacy network infrastructure over many years leveraging seasoned support personnel. With increased competition, speed to market is critical for survival and growth in the communication business. This means that the rapid rate of change in networks is limiting the service provider to adequately train and ensure experienced personnel to help in the operations and maintenance. In the short term, although more resources can be put on the introduction of the network/service, providers need a new support paradigm to reach a sustained cost competitiveness.

Network-level support

An emerging maintenance model enhances the current product support model for the Service Provider in two dimensions: 1) provides network-level support of the infrastructure, supplemented by product-level support, and 2) provides a single point of contact to work issues across multiple vendors. Leveraging a network integrator as the single point of contact, providers no longer have to chase multiple vendors for answers or to resolve disputes.

A single-point-of-contact network maintenance model enables service providers to realize lower operating expenses and allows them to focus on service delivery for revenue growth. They no longer have to dedicate personnel to different vendor interfaces; instead, they are free to use these resources to build out their next-generation networks and rolling out the new IP-based services that their customers demand.

As an evolution of product-level support, the network support model continues to provide the Service Providers support-level metrics based on the acceptable respond, restore and resolve times. It is also similar in the operational processes between the Service Provider and the network integrator. Service starts with the service provider identifying a network or service problem and initiating a request for help. The network integrator will require all the relevant network information (e.g, management logs, alarms and traps, and other historical data) to analyze the problem with in-house experts and to isolate the network fault. The network integrator then contacts the service provider, identifies the root cause of the problem, works with the third-party vendors on behalf of the Service Provider to get the problems fixed, and recommends the appropriate corrective resolution back to the service provider.

The network integrator's role doesn't have to be limited to fixing problems as they occur. This model brings further value to the service provider where the network integrator can also serve as a key technical consultant to help to create an end-to-end network-level map and identify potential "hot spots" of potential future issues. Familiar with the current network, the network integrator can also help assess the provider's roll-out plans, pointing out where configuration issues or capacity limitations could lead to migration problems.

Under this model the network integrator becomes a key resource for other network-related issues. For example, a service provider can contact the network integrator for its expert advice about network configuration. The service provider presents what it's planning to do; the network integrator assesses issues like configuration and capacity, as well as susceptibility to outages. Essentially, the network integrator has become a trusted partner to the service provider, sharing insights and expert advice from a global experience of multiple customers and applying that knowledge to the Service Provider's unique network and service plan to help increase efficiency, speed service rollouts, and boost revenue.

Comparison of product, network and service-level support

KEY ATTRIBUTES	PRODUCT-LEVEL SUPPORT	NETWORK-LEVEL SUPPORT	SERVICE-LEVEL SUPPORT
Support trigger	Service provider	Service provider	Both service provider and solution integrator
Network trouble analysis	Service provider	Network integrator via remote login	Network Integrator using automated, intelligent tools
Fault Isolation to product-level	Service provider	Network integrator	Network integrator
Performance metrics	Product support-level metrics (respond, restore, and resolve times)	Network support-level metrics (respond, restore, and resolve times)	Network Support-level metrics and end-to-end service-level KPIs/KQIs
Critical success factors	Product experts	Solution experts	Solution experts, predictive modeling, and solution knowledge base

Service-level support

In today's world, the service providers need a type of support to match the evolution of a network performance model to a service performance model that operates on the converged network. A service provider purchases an end-to-end solution that promises to deliver a range of multimedia, real-time services on a shared IP infrastructure over various access technologies to reach the subscribers. The most important maintenance question the service provider can ask is: Am I optimizing my revenue? The answer depends on whether subscribers are successfully receiving service with the quality they expect. Poor service quality results in customer loss/churn, and limits the repeated

service growth of revenue-generating applications with existing and new customers. A solution integrator that looks only at the product and network performance will not be able to answer the question.

Thus, the customer service support model must include the services (or applications) as well as the network. Service providers need the solution integrator to accept assistance requests that are not necessarily tied to any network problem. Services may be unsatisfactory even if the network domain appears to be running smoothly. Conversely, network problems may or may not be affecting end services. Integrators must accept that this is part of the maintenance support problem.

With service-level support, the integrator and service provider work in partnership to maintain the services offered. The service provider must still control the operational aspects of its services — provisioning, billing, fault, and performance management — and all the other functions typically addressed in a network or services operations center. Meanwhile, the integrator provides the experience, tools and models to support the service provider and helps it to rapidly sectionalize, prioritize, and resolve service-affecting issues. It is this new synthesis, between the service provider and network integrator that will allow delivery of innovative end-user services with the quality and availability that subscribers demand.

To achieve the level of unified support to a service provider, the integrator will need access to the data the service provider gathers. The network integrator structures, correlates and mines that data so the data can be used as meaningful information. That means deriving the critical statistics and displaying them onscreen, giving the network integrator a window into service performance, trends, and problems in the making. Knowing how to calculate the right KPIs from the mass of available data is a very valuable asset for the network integrator. Combining the KPIs with the expertise to diagnose and **prevent** problems is the key ingredient in optimizing service availability for a provider. It is the insight that a network integrator gains after many successful installations and support of many networks around the world.

The DNA of a proactive service-level support model

In addition to the elements of product and network-level support, there are four essential ingredients to the Service-Level support delivery methodology:

Predictive modeling

The predictive model specifies the relationship between the performance of network elements and the service performance experienced by end-users. This mathematical model is a representation of the solution architecture and end-user service delivery. The predictive model is crucial; it enables the network integrator to closely estimate equipment performance values, plus the quality and availability of the service delivered to the end-user. With this information, the network integrator can recommend changes to the service provider that could improve both indicators — to delight end-users on the services they are purchasing.

Near real-time monitoring

When a network integrator collects information from all the levels of the deployed solution, it can determine in near real-time that there are maintenance issues that need to be addressed — while they are still potential rather than actual problems. In this way the network integrator supplements a service provider's network operations center (NOC).

Proactively managing the network is a must for avoiding quality of service degradation. To stay ahead, network performance has to be continuously monitored to reveal which resources are at risk. Thus potential problems can be addressed before they impact end-users or cause delays and service outages.

A network integrator who delivers real-time monitoring leverages the same data available to the NOC. However, the network integrator works with the NOC to apply the experience and expertise gained from all supported customers globally to speed its maintenance efforts.

CASE STUDY: AN IPTV OPERATOR COULD REDUCE FAULT ISOLATION TIME FROM 4 HOURS TO MINUTES.

Problem

A failure occurred on a Network Interface Card (NIC) in an A-server. Normally, this event would have no impact on end-to-end service, because the NIC would fail over to a backup card and service would not be interrupted. In this case, the failover was successful so the fault management staff believed the problem was minor. However, the backup board was connected using a bad cable which resulted in the board negotiating a much lower interface speed to the network. As a result, packets began to be dropped and video quality was affected for all the channels served by that A-server.

Relying on only the traditional maintenance support, this problem took the service provider a significant amount of time (approximately 4 hours) to understand, diagnose and resolve. When video quality problems are reported, maintenance staff begins at the set top box and works back into the network looking for problems. It was logical to examine the D-server because it is supposed to fill in any lost packets in the original stream to the STB. This requires logging into multiple places in the network, running commands and looking manually at the results.

Solution and benefits

With the service-level support coupled with near real-time monitoring, the maintenance staff has access to the data necessary to understand the problem, without manually logging into the network. The maintenance engineers will be able to immediately point to the A-server as the source of the problem, and compares the configuration of the NIC speed with the expected value, creating an alarm on the discrepancy. This alarm includes a recommendation to check the cable, and the problem is resolved within minutes instead of hours.

Of course, this NIC speed alarm can be created even before the failover happens! In this way, maintenance staff can alert the service provider of the problem before it ever affects subscriber video quality.

Proactive response

Thanks to predictive modeling and near real-time monitoring, it's possible for a network integrator to be truly proactive, eliminating potential trouble spots and other possible problems before they can become real threats. For example, if the monitored data shows the KPIs disagree with the assumptions in the predictive model, the network integrator can project the impact on service users and alert the service provider with a suggested course of action.

CASE STUDY: NEAR REAL-TIME MONITORING ENABLED A WIRELESS CARRIER TO GAIN OPERATIONAL EFFICIENCIES FROM INCREASED VISIBILITY TO KPIS/KQIS

Problem

A wireless carrier has no consistent visibility into their regional and national operational performance levels. The conventional product-level maintenance does not provide the visibility that the carrier needed to measure and analyze the performance of network operations across various markets to improve operational productivity and to assure its customers the highest level of performance and availability

Solution and benefits

The network integrator and the carrier first jointly determined the critical KPIs and KQIs that enable the provider's managers and executives to set, measure and achieve operational objectives for their voice, data and next generation wireless networks.

The network integrator then delivers a service-level support service which provides near real-time monitoring of the KPI and KQI metrics to enable the provider to identify potential problems and take corrective actions before the problems occur.

The information also allowed the provider to take proactive actions to gain operational efficiencies, which resulted in operating expense savings from more effective network maintenance.

Knowledge base

A knowledge base is an expert system with highly interactive database that makes all of a network integrator's experience instantly accessible to its support engineers anywhere in the world. The knowledge base includes everything from vendor patches and deployment details to end-to-end problems — and their solutions. The knowledge base allows a network integrator to access the details of numerous successful deployments when helping the service providers to resolve service-level issues.

The knowledge base also drives proactive changes to customer configurations based on other customer deployments.

Conclusion

Rising consumer demand for innovative services is driving the transition to next-generation networks. But rolling out services is not enough. Consumers expect — and often demand — top quality across all services anytime, anywhere, on any device. Service providers must meet customers' demands or risk losing them to the competition.

Consumer demand also is driving the transition to next-generation maintenance models. In the first stage, service providers and network integrators will begin to work as partners. Together, they will optimize the end-to-end availability of the network. This is a more efficient maintenance model; providers will be able to offload much of the hard work that constitutes network maintenance, focusing instead on services and service quality. In the second stage, network integrators will take on application-level maintenance. Through predictive modeling, service monitoring, and proactive responses, this maintenance methodology can optimize service availability. At that point service providers can concentrate on delivering innovative services, confident in their superior quality of service, and attracting customers from their competition.

About the authors

Peter Peng
Director, Business Development, Services Business Group, Alcatel-Lucent

Peter has over 20 years of experience in managing, developing, and marketing telecommunications solutions for the global service providers. Peter is currently leading the development of next-generation solutions serving global carriers' multi-vendor network support needs.

Prior to his current role, he was Director of Application Planning & Portfolio Management, responsible for network planning & design, technical marketing, and portfolio planning of optical networking products. In this capacity, he advised numerous telecommunications operators world-wide in the development of long-haul, metropolitan and access networks.

Peter began his career as an engineer in Bell Labs, pioneering the development of Fiber-To-The-Home systems and designing large scale Digital Cross-Connect Systems. During the early 1990s, he and a team of engineers at Bell Labs successfully developed the first photonic cross-connect system. Peter holds three U.S. patents.

Peter received his BE degree in Mechanical Engineering from Cooper Union School of Engineering and his MS degree from California Institute of Technology.

Dr. Richard Jordan
Solution Architect, Services Business Group, Alcatel-Lucent

Richard is the technical leader for development of Alcatel-Lucent's End-to-End Service Availability Support offer for IPTV, IMS, and UMTS network technologies.

During Richard's 22 years as a scientist in the telecommunications industry, he has led many major projects from concept to delivery. He is an expert in the optimization of next-generation network technologies, OSS architecture, and network management solutions.

Richard has held positions in systems engineering, software development, architecture, product development, technical sales, and product management at AT&T Bell Laboratories, Lucent Technologies, and Alcatel-Lucent. He has worked at customer sites and consulted with service providers around the world. Richard holds two patents on network traffic management solutions.

Richard has an MA in Mathematics from Oxford University, the Certificate of Advanced Study in Mathematics from Cambridge University, and a PhD in Theoretical Physics from the University of Texas.

Robin Plessl
Director, Product Management, Services Business Group, Alcatel-Lucent

Robin is a leader in the Alcatel-Lucent Technical Support Product Management team with responsibilities that include market launch of the End to End Network Support Service for Next Generation Networks and Multi-vendor support. She has over 25 years experience in the communications industry, supporting both enterprise and carrier customer markets. Her specific areas of expertise are research & development, engineering, solution/product management, supply chain, and sales.

Robin received her BS degree in Chemistry from Lebanon Valley College and her MBA from Lehigh University.

Alcatel·Lucent