

Network Performance Infrastructures: Not Your Father's NMS



Apparent  **Networks**
Network Intelligence for Immediate Solutions



Co-written by:

Prasad Calyam, Internet2/OARnet



Internet2 is a consortium being led by 205 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies, accelerating the creation of tomorrow's Internet. Internet2 is recreating the partnership among academia, industry and government that fostered today's Internet in its infancy. The primary goals of Internet2 are to:

- Create a leading edge network capability for the national research community
- Enable revolutionary Internet applications
- Ensure the rapid transfer of new network services and applications to the broader Internet community.

OARnet, an Ohio-based Internet Service Provider, has been helping its clients reach their educational and business goals on the Internet with a variety of support services since 1987. More than just another ISP, the leadership at OARnet, intent on building the organization and its network to an unsurpassed level of superiority, has risen to a new level: Internet Solutions Partner.

Loki Jorgenson, Apparent Networks

Apparent Networks

Apparent Networks is a leading innovator of network intelligence software. Apparent Networks' technology, AppareNet, a network intelligence system, operates non-intrusively on live networks, to and from any location worldwide. Without requiring specialized hardware or remote agents, AppareNet views the network from the application's perspective. In doing so, AppareNet rapidly identifies the locations and causes of network bottlenecks anywhere in the world so that companies can boost the performance of, and gain more value from, the network infrastructure they already have. Apparent Networks improves its customers' businesses by helping organizations reduce operational costs, increase IP availability, and protect revenues.

This report in whole or in part may not be duplicated, reproduced, stored or retransmitted without prior written permission of Apparent Networks. All opinions and estimates herein constitute our judgment as of this date and are subject to change without notice. Any product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Introduction

Network Management Systems (NMS) were last decade's answer to the ever-increasing demands of networks. They offered a best-effort solution that promised to make sense of the morass of data extracted from multitudes of devices. However, the assumptions behind their design and implementation rarely apply any more.

The old client-server approach provided a relatively controlled environment where well-defined processes took place. The goal of IT was to increase the productivity of specific business processes and the budgets were relatively flexible. The issues facing network engineers revolved primarily around management and users had only limited contact with the network.

Today, application performance defines the success of networks, and almost every aspect of business depends on applications that depend on the networks. The networks grow and change more rapidly, with critical parts outsourced to ISPs and other providers. Information Technology has become as accountable as any other business unit to show ROI and even generate revenue.

Today management of your networks just isn't enough. Application assurance is critical and this requires a new approach. Cumbersome, high maintenance Network Management Systems have begun to give way to rapidly deployable Network Performance Infrastructures (NPI) that support access for all stakeholders in the network community -from the experienced network engineer, to support staff and helpdesk and, most importantly, to the end user.

By definition, a well-implemented Network Performance Infrastructure provides a high level of immediate feedback from the network so that all forms of application performance can be assessed and guaranteed, and problems can be proactively identified and resolved, leveraging the advantage of effective participation by all stakeholders.

Worlds Apart

The transition from early data networks to modern day, high performance communications systems has changed the playing field dramatically. Very little is recognizable from the origins of networking.

In the old world:

- you owned or controlled most of the networks your key applications depended on;
- you could predict where critical traffic would flow and when;
- complex, slow deployment, agent-based systems were approved and implemented because the problems client/server created were new, acute and difficult to resolve;
- you dealt with a dog's breakfast of vendor-specific protocols that often forced you to maintain relatively homogeneous systems with clear functional boundaries;
- you had smart engineers designing and managing the networks directly;
- your users and customers had limited expectations of the networks - down time was expected and acceptable;
- business processes that depended on your network were limited and well-defined;
- IT was implemented to generate productivity increases and budgets were generous.

In today's world:

- your business depends on networks you do not own or control (ISP, ASP, customer, supplier, etc.);
- you can't predict where tomorrow's traffic will flow or what will break next;
- there is less (or even no) time and money for deployment or maintenance of big, complex network management systems;
- it's "IP everything everywhere" with legacy systems in between;
- your users and customers have higher expectations and your business network is critical – even brief outages can be disastrous;
- you rely increasingly on helpdesk and support staff with limited network expertise and few tools to do their job well;
- networks are mission critical - everything seems to depend on them;
- IT is expected to align with other business practices, including cost reduction and even revenue generation.

There is little wonder that the Network Management Systems of the last 20 years do not address today's needs. What should you be looking for instead?

Network Life Cycle

Today's networks are highly dynamic- they evolve within a never-ending cycle of planning, deployment, maintenance, and upgrading. The traditional NMS has serviced only a very small part of the Network Life Cycle (NLC) and typically at a very high cost.

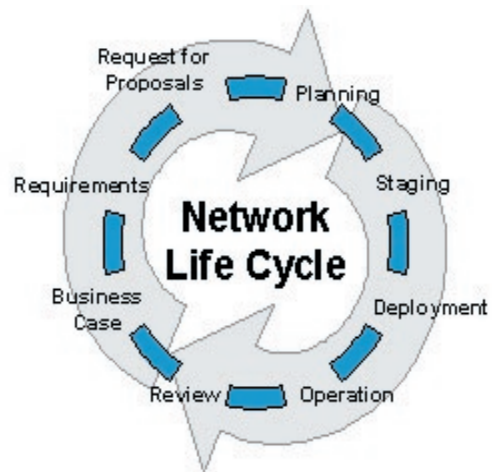


Figure 1 – The Network Life Cycle defines the essential processes of the network.

Operation is still the obvious part of the process but the bounds of today's network operation have expanded dramatically to include applications, data centers, outsourced resources, and helpdesk/support for both internal and external users. NMS have primarily focused on data gathering and device management. This is useful but limited to networks that you own and have control over.

Incremental improvements in various management technologies have armed network engineers with increasingly sophisticated measures of network performance, including addressing aspects of various forms of application performance. They provide a view of performance defined in terms of a particular application or use of the network.

For example, consider the wildly differing requirements of data storage and Voice-over-IP. Each tool represents an example of a Network Measurement Infrastructure (NMI).

NMIs provide a coherent, contextual basis from which to make timely and specific measures particular to a use, application, or environment.

However, the performance of the network is not fully defined by any one of those particular views. These new technologies are highly attractive to anguished network personnel and yet they do not offer a sufficiently complete solution to warrant investment in time and money.




Figure 2 – The data pyramid on the left depicts the levels of refinement that transform massive amounts of data into various measures (information) and finally into intelligence. The inverted Infrastructure pyramid on the right shows the corresponding increase in value for each type of infrastructure.

A fully coherent NPI offers an integrated view of the network based on all aspects of the end-to-end network path. This view includes perspectives from various NMI-type measures and assessments, and it also includes an integrated analysis of the overall performance. Further, an NPI provides that view relative to the applications that use the network, as well as the members of the network community (users, helpdesk, application support, engineers).

Approaching the Mountain From All Sides

Although NMIs are increasingly available and well-developed, they tend to be narrowly focused and relatively inflexible. The focus on application performance is driving the emergence of new metrics and measurement technologies associated with them. The need for a coherent, overall solution to network performance is driving the transformation of NMIs into a broadly accessible and flexible NPI:

- o Seeing the application’s view, end-to-end, includes all the components that impact performance;
- o Deploying the requisite infrastructure rapidly, and on-demand, ensures that scarce resources are applied when and where needed;
- o Being able to see into networks you don’t own means you can out-source with confidence;
- o 24/7 monitoring ensures that “network awareness” gives you immediate feedback;
- o Real-time responsiveness based on current conditions reduces the dependency on historical data and makes the infrastructure more adaptable;
- o Emphasizing “effective” over “absolute” means that you resolve the most common and most expensive problems most quickly;
- o Distinguishing clearly between TCP/IP and higher up the stack means that problems with applications can easily be separated from problems with the network;
- o Your network’s needs should be reflected in your performance infrastructure instead of fitting your network to the infrastructure;
- o Capitalizing on every aspect of your existing infrastructure, including the effective participation of your end users, means nothing goes to waste;
- o Providing appropriate access to all levels of management and the user base means that everyone can usefully participate in the process.

When you start to put these principles to work, a meaningful picture starts to emerge. Let's consider a couple of examples, one each of an NMI and an NPI.

Use Cases and Exemplary Technologies

A use case for an NMI is found within Internet2 where ambitious video-based projects are being deployed. The Video Commons project aims to facilitate the interaction of researchers and students at universities connected to Internet2. End-users cannot be expected to deal with network issues directly, so tools are being developed that enable the end-user (and network support) to see end-to-end, helping them identify and resolve their own video-related problems.

An NMI typically assesses the performance for a specific application, stakeholder, and/or environment. In this case, it is specific to video using the H.323 protocol, primarily supports the end-user, and is directed at the WAN environment.

NMI Example: H.323 Beacon

The **H.323 Beacon** is one of the new-generation network measurement tools which knows about the application itself and can troubleshoot end-to-end application performance problems specific to its character. The H.323 Beacon emulates H.323 Videoconferencing applications and possesses the diagnostic capability to detect firewalls or NATs along the path that hinder call establishment and audio/video media exchange.

It can be used to measure, monitor and qualify the performance of an H.323 Videoconference session. It can help an end-user, or network engineer, or conference operator resolve issues prior to, during and even after, videoconferencing sessions. It provides H.323 protocol specific evidence, and other information necessary to troubleshoot H.323 application performance problems in the network and at the host (end-to-end).

H.323 Beacon uses a distributed client/server architecture - the client actually refers to an end-node and the server can be visualized as a core-node. Testing between end-nodes can be achieved by using a number of core-nodes along a test path. This architecture facilitates the H.323 Beacon to perform end-to-end measurements related to H.323 videoconferencing sessions.

About H.323

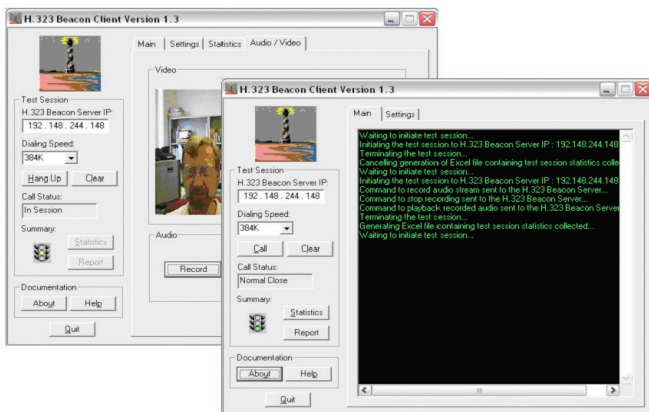
H.323 is an umbrella standard from the International Telecommunications Union (ITU) for developing real-time multimedia communication applications such as audio/videoconferencing on packet-switched networks (Internet). The popularity of the H.323 technologies has been demonstrated by the billions of minutes of voice and video traffic seen on the Internet every month.

H.323 is a recent technology; thus, there are few affordable and appropriate diagnostic tools available to the end-user or network administrator to identify and troubleshoot performance problems related to H.323 applications.

The H.323 Beacon has three levels of test status: "In Session", "Normal Close" and "Exception Close". The "In Session" and "Normal Close" status depict the user initiating the test session and the user ending the test session. The "Exception Close" status depicts a potential performance problem either at the beginning or during the test session that caused the test session to close. The "Exception Close" status is always accompanied by an "Alarm" that indicates the possible cause of the performance problem. A few examples of the alarms are: "No Internet connectivity", "Possible firewall/NAT presence obstruction", "Excessive network congestion", "Insufficient bandwidth" and "Signaling incompatibility".

The H.323 Beacon also collects local network statistics such as frame rate, delay, jitter and packet loss in real-time during the test session. Graphs for protocol signaling and network characteristics are generated with colored watermarks for good (green color), acceptable (amber color) and unacceptable (red color) values of delay, jitter and packet loss.

With the audio / video loop back feature, the "local" audio/video quality as experienced at the remote end can be seen locally. Local video can be recorded at the Beacon server in AVI, MPEG, or QuickTime formats. Local audio can be recorded at the Beacon server in WAV format. The recorded audio/video files can be played back at the client by FTP-ing transparently from the end-users perspective.



At the end of the test session, a detailed test session report is generated with information regarding exceptions, overall network performance values and forward-reverse traceroutes that could potentially be handed over to more experienced support staff to resolve the performance problems experienced by an end-user. The report feature is a relief to support personnel who are faced with situations where their users complaints are in the form of "I experience bad video during my videoconferences!" instead of providing them adequate evidence to troubleshoot videoconferencing problems.



With the intuitive GUI, the added ability to conduct client-to-client, client-to-server and server-to-server H.323 application performance tests, along with the ability to customize protocol specific parameters for tests, the Beacon is a comprehensive solution to troubleshoot performance problems before, during or after an H.323 Videoconference.

H.323 Beacon Use Case: Educational Classroom End-user

A university's IT Manager purchases expensive videoconferencing equipment to facilitate Distance Learning at his campus. When a Professor wanted to deliver a lecture to a remote class via videoconferencing, the IT Manager was called in because the Professor was not able to connect to the remote classroom's videoconferencing equipment. The IT Manager calls up the IT staff at the remote classroom site, and is advised by the remote site support person to run a H.323 Beacon test to the remote classroom location. The H.323 Beacon reports a firewall in the Professor's campus LAN that is configured to block ports required for H.323 protocol signaling. Once the firewall port blocks were eliminated, the Professor was able to connect to the remote classroom equipment, although both sites observed bad video and audio reception. Simple ping tests did not indicate a performance problem. After running the H.323 Beacon again, they obtained the graph plots of the network statistics. The jitter and packet loss curves had crossed critical thresholds which represented "Unacceptable" values for a high quality videoconference session. Upon running the H.323 Beacon tests through available H.323 Beacon Servers along the path between the remote classroom and the Professor's location they realized that one of the intermediate links was experiencing significant packet loss. Upon resolving the link problem, the Professor and the remote classroom were able to successfully collaborate using videoconferencing.

NMI → NPI

The advantage of the NMI can be critical to the success of the application it is designed for. In the case of the H.323 Beacon, users now have visibility that supports them to respond to a particular set of problems.

Extending this capability and integrating it across a range of other measures/analyses brings this approach closer to the desired NPI. As a comparison, the AppareNet software offers a view of the network that is generic and applicable to all applications. And, by employing an expert system, AppareNet leverages its information into context-specific intelligence that can support a user or a network engineer. Within the AppareNet Triage offering, it subsequently bridges across several of the key obstacles to effective network performance management.

NPI Example: AppareNet Triage

AppareNet Triage answers the requirements of Best Practices by integrating the benefits of a comprehensive NMI technology into an NPI offering. The core AppareNet technology identifies Layer 3 as the foundation of all application performance and performs measures/analyses that identify and isolate key features of the network path in real-time. By addressing the issues of rapid on-demand deployment, remote data gathering, continuous monitoring, user-specific views, and universal accessibility, the AppareNet Triage solution takes the NMI approach of the core product the next step.

The AppareNet measurement technology follows a point-and-shoot methodology, requiring very little pre-configuration to execute a comprehensive analysis of almost any network path. A network engineer might typically only select the deployed point of view from which to measure and enter the IP address of any arbitrary end-point.

Then, simply by pressing “Start”, the NIS discovers the end-to-end path, identifies all visible Layer 3 elements, and begins comprehensive testing without further prior knowledge.

Hop	Severity	IP Address	Host Name	Packets Sent	Packets Lost	Packet Loss (%)	Bandwidth (2way, Mb)
5	OK	66.59.190.10	ge4-0-peera-stlwa.ip.grouptelecom.net	1350	0	0	
6	OK	129.250.10.33	pt4-5-3-0-r03.stlwa01.us.bb.verio.net	1350	0	0	
7	OK	204.255.169.25	POS2-0-BR1-SEA1.ALTER.NET	1350	0	0	
8	OK	152.63.105.22	0-so-7-3-0-r02.sea1.alter.net	1350	0	0	
9	OK	152.63.2.134	0-so-7-0-0-r02.sea1.alter.net	1350	0	0	
10	OK	152.63.2.37	0-so-4-0-0-r02.atl4.alter.net	1350	0	0	
11	OK	152.63.89.233	0-so-3-0-0-r02.dem4.alter.net	1350	0	0	
12	OK	152.63.89.253	178.atm7-0-gw7.dem4.alter.net	1350	0	0	
13	OK	157.130.164.86	lucent-co-gw.customer.alter.net	1350	0	0	
14	Diagnostic	192.11.229.2	2.229.11.192.in-addr.arpa	1350	336	25	

Severity	Diagnostic	Certainty (%)
Diagnostic	Half/Full-duplex conflict detected	82

Severity	Observation	Frequency (%)
Warning	Packet loss detected [Exceeds reasonable tolerances for this network path]	25
Warning	Packet reordering detected	7
Info	Insufficient statistics to generate all measurements [Packet loss may be interfering with test results]	-

The real-time result is an end-to-end assessment of the network’s characteristics in detail, summarized simply and effectively as either optimal (green light) or dysfunctional (red light). In addition, the expert diagnostics system can identify a wide variety of common problem behaviors and will tell the user what the cause is and where to look.

AppareNet Triage combines this measurement and diagnostics capability with several other functionalities:

- o It can shoot into and through unknown networks;
- o The deployed network sampling component (the point of origin) can be deployed and re-deployed on-demand anywhere, even behind firewalls;
- o For rapid, one-time deployments to resolve remote issues, a throw-away version of the sampling component can be delivered by e-mail or the Web for manual deployment;
- o Simplified interfaces with "smart" views support end-users to transparently test their own connections without consultation with network;
- o The intelligence gathered by end-users travels with their trouble calls so that network engineers can see what the user saw;
- o It can monitor critical paths on an on-going basis.

This approach ensures that the NPI can be rapidly deployed into any network as needed, and allows all members of the network community to contribute to the rapid identification and resolution of issues impacting application performance. The solution can be flexibly applied to the needs of Enterprise networks, ISPs and ASPs, and Network-Dependent Vendors.

AppareNet Triage Use Case: Network Dependent Vendor

A customer purchases a product (i.e. CRM system, ERP system, storage servers) from the vendor. At some point in their experience with the product, the Customer encounters a performance problem. The Customer contacts the vendor's Customer Support line for assistance. Customer Support asks the standard questions and cannot quickly diagnose whether or not it is the product or the customer's network that is causing the problem. Using AppareNet, Customer Support then executes a network test on the customer's remote network. In less than 5 minutes, the AppareNet test results will identify if there is a network problem or not. If the test results do uncover a network problem, the exact cause and location will be pinpointed. Customer Support can then route this information to the remote customer for resolution on their end. With AppareNet, enterprise hardware or software vendors can drastically reduce the mean time to resolve customer support issues while increasing the value of their customer support organization beyond their basic product offering for improved customer satisfaction.

As AppareNet's NPI offering is expanded, additional metrics and application-specific views will be integrated to potentially include the metrics and analyses associated with NMs like H.323 Beacon.

Best Practices for an Effective NPI

The case for a modern Network Performance Infrastructure can now be made. The hallmarks of a complete solution are defined by Best Practices.

Here are the keys to look for:

1. **Continuous monitoring of performance** (not just availability) as an essential starting point
 - 1.1. ideally a Layer 3 or 4 demarcation point separates network issues from application
 - 1.2. application performance is assessed specific to each application type
 - 1.3. "smart filtering" that limits or eliminates redundant or false positive notification
2. **Rapid response** to performance problems that slip through the cracks
 - 2.1. a real-time measurement/assessment/problem diagnosis capability
 - 2.2. automated expert analysis
3. **Rapid deployment** and configuration
 - 3.1. capable of on-demand responsiveness
 - 3.2. can be deployed remotely within otherwise inaccessible networks
 - 3.3. no need for a priori knowledge of the network
 - 3.4. auto-discovery and self-configuration
4. **Empowerment through access** to the performance infrastructure throughout the Network Vertical
 - 4.1. network engineers
 - 4.2. helpdesk and support engineers
 - 4.3. end-users and customers
 - 4.4. strong share-ability of results between all the members
 - 4.5. intelligent support for all members
5. **"Smart"**, not dumbed-down, client-side views
 - 5.1. application-specific analyses and diagnostics
 - 5.2. action-oriented expert feedback
6. **Extensibility** in the form of adaptability and configurability
 - 6.1. in application type and requirements
 - 6.2. vertical member role and responsibilities
 - 6.3. variable data source and performance analysis methodology
 - 6.4. follows work process
7. **Intelligence archiving** that supports sharing of network data between organizations
 - 7.1. Methodologies that provide visibility into networks that you don't own
 - 7.2. Federated access and data sharing
 - 7.3. Anonymization and secured data access

The Future of Utility

The elements of the complete Network Performance Infrastructure are not just for today's needs. They are congruent with the trends for near-future technologies. Utility and on-demand computing will require the same capabilities delivered in the same dynamic and adaptable package. Instead of serving the user and their associated support desk, the NPI will be the source of direction and performance analysis for automated and adaptive computing resources. To date, users have been the "performance monitors", limited in their ability to evolve the process from using the phone to call helpdesks. The lack of effectiveness in this approach has driven the trend to NMIs and subsequently NPIs. However, soon we can expect the user to be relieved of the tiresome task of monitoring their own environment. As NPIs provide the basis for self-healing and adaptive networks, applications will enable the optimization of their own performance automatically.

References

- o H.323 Beacon
<http://www.itecoho.org/beacon/>
- o Apparent Networks white papers:
<http://www.apparentnetworks.com/wp/>