

CommNews January 05 Article

Vendor contact (and byline):

Dave Trowbridge
Director of Marketing
Network Physics
davet@networkphysics.com
650-230-0970
650-230-0909 (fax)
650-714-1143 (cell)

Vendor information:

Network Physics
491 Fairchild Drive
Mountain View, CA 94043
650-230-0900
650-230-0909 (fax)
www.networkphysics.com

The central function of IT in the modern enterprise is to assure the performance, integrity and security of the application infrastructure. Network managers have struggled for years with their part of this effort, piecing together increasingly unwieldy collections of point solutions and largely device-oriented data sources reporting to a monolithic, central repository in an attempt to “defend the network.” Nonetheless, too often, all lights are green but users still complain about poor application response, and all the other management teams point at networking. Network application management addresses this problem by focusing on the network’s role in infrastructure assurance while also giving network managers cross-domain insight into server and application performance and utilization. The result is better coordination between teams, leading to a more efficient management of the application infrastructure.

However, a traditional, centralized, largely device-based approach is not up to the four challenges of network application management in a global enterprise:

- Real-time coverage
- A unified logical and physical view of the network
- Scalability
- Management agility

What is required is a distributed system that moves much of its management intelligence to the edge of the network, mirroring the “dumb network” model that has made TCP/IP networks so successful. As we shall see, this is most easily accomplished with a new technology category: flow-based data collection and analysis, as found in products from vendors such as Compuware, NetScout, Network Physics, and Packeteer.

Real-Time Coverage

A solution that attempts to assemble a coherent view of network operations and their impact on application performance by aggregating data from thousands of network elements, agents, and probes into a master database cannot practically provide real-time coverage of a global network. The network impact of such management traffic is enormous; indeed, some large financial firms have built parallel networks just for management data. The database impact is similar. But worst, perhaps, is the burden of managing thousands of data sources. As a result of these practical limitations, RFIs for centralized systems generally specify 10-15 minute averages; perhaps 5 for those companies that can afford it. And this not only skews averages away from the true traffic peaks, but misses many transient events—the most difficult to troubleshoot.

Flow-based data, on the other hand, can be gathered by aggregating the data from TCP/IP packet headers as they pass through a data aggregation point, such as a spanning port on a switch, so it is not necessary to manage thousands of sources. By itself, this flow data supplies detailed information on performance and utilization; other information, from BGP and traceroute, can be correlated with the associated flows to integrate route and ISP/AS data as well to deliver a range of metrics usually requiring four or five point solutions. This flow-correlated data, which combines Layer-3 and 4 information, is also more compact, since their performance and utilization information is independent of the number of devices supporting the flow.

Even so, a centralized, flow-based solution would eventually be overwhelmed by the same problems. But a distributed solution moves most of the system intelligence to the edge, so that only relevant information is moved to and analyzed by the central console. By “moving knowledge, not data,” a distributed system greatly reduces both the network load and the necessary size of the central database, permitting real-time operation. A federated database design can assure transparent access from the central console to the more detailed data stored on the intelligent edge devices as necessary.

Unified Logical and Physical View

Unifying a logical, business-centric view of the network with a physical view is critical for efficient troubleshooting. A logical view of the network presents data in the context of applications and the user experience, which is how performance is perceived, reported and judged, and of the business entities involved (e.g., branch offices, departments, buildings, users, server clusters, business applications, and end-to-end business services), which permits rapid problem prioritization. (It also supports functions such as planning, baselining and chargeback). The physical view is needed to support drill-down from the symptoms, which are reported in business terms, to the devices responsible.

A distributed solution can more easily maintain these contexts of network data at all scales, from a global overview to details of individual IP conversations. It can do so by using intelligent aggregation (grouping) and “topping” (selection of N top groups by a chosen metric) algorithms to create digests for the global management console, as well as the federated data storage to make possible the preservation of data details on the edge appliances that would otherwise be lost in aggregation and statistical synthesis. To avoid

network overload, a centralized system must top locally, then group on the central console, which alone has the intelligence to do so. Averages are thus based on already topped data selected without grouping intelligence, and data has been lost.

It is also difficult to assemble a logical view of the network from device-level data such as SNMP and RMON-2. This requires hand-coded “maps” of the relationship between a set of devices and the business processes and entities that they support, which are brittle in the face of infrastructure change, requiring reprogramming. Some automation of this is possible, at the cost of much greater management complexity. The processing and database load are also high, especially in a centralized system.

By contrast, a multi-dimensional, logical view of the network is implicit in a flow: an end-to-end IP conversation (most frequently TCP or UDP), across any IP network type, between two business entities, a resource provider and a resource user, whose conversation or transaction supports a business relationship (e.g., application and user, server and client). With appropriate aggregation (grouping), a flow-correlated solution delivers application, user and business context unaffected by changes in the underlying physical infrastructure while maintaining drill-downs to the physical devices involved. Integration with a Layer-2 solution (e.g., SNMP) permits even further drill-down, into the internals of specific devices.

Scalability

A distributed solution avoids scalability issues the same way that TCP/IP does: by pushing intelligence—data collection, analysis, storage and, most important, action—to the edge. (One might call it “managing a grid with a grid.”) As we shall see below, a distributed system’s greatest contribution to scalability is its “management agility.”

Device-based systems are inherently difficult to scale, especially given the autonomous nature of TCP/IP networks, which manage much of their own flow rate and routing. The network and database load, and above all the management burden represented by thousands of devices, agents and probes reporting to a central console are necessarily high, and there is no economy of scale—these problems just get worse the larger the network.

Thus, flow-correlated technology does its part here: its metrics are largely infrastructure independent, and millions of flows can be monitored from a single point, rather than depending on the complexity of multiple SNMP, RMON and application response agents. This also makes a flow-correlated solution easier to distribute than a device-based solution.

Management Agility

The goal of enterprise IT is to adapt the application infrastructure to business processes, not vice versa. That’s just as true for the process of infrastructure assurance as any other: solutions should adapt to the management structure of enterprise IT. A network application management solution should adapt easily to support the most efficient

division of managerial oversight, e.g. global/regional, server/application/network or manager/C-level. It should also let users explore management data however they wish.

This is inherently difficult in a centralized system. Most tasks can and should be handled at the local level. A centralized system makes this difficult, however, since the data needed for troubleshooting and planning can only be supplied by a central analysis engine. Global managers can become overwhelmed by local problems that don't need their attention, while local managers have to refer to a central database, requiring complex role and permission setting and often delivering less detailed data than would have been available locally, due to the necessity of "topping" data to preserve storage space in the central repository.

By contrast, in a properly-designed distributed system (i.e., with a federated database), local managers enjoy immediate access to the data stored in the local, intelligent devices, while global managers see only digests or summaries until they need more detailed information. This also allows the enterprise to preserve its chosen balance between regional and global management. Furthermore, integration at the local level with smaller, more manageable SNMP-based solutions, gives the Layer-2 drill-down to the managers with day-to-day responsibility for the devices involved.

ILLUSTRATION HERE

The greater ease with which flow-correlated technology supplies application, user and business context also contributes to management agility. In contrast to a device-based approach, a flow-based solution can rapidly identify a problem as server, application or network-related, allowing the right team to get on it quickly and avoiding any finger-pointing.

Finally, good reporting is a necessity. Network application management requires reporting that scales from real time, for troubleshooting and triage by on-the-spot managers, to long-term, for tasks such as planning, baselining and chargeback, and higher management levels. Centralized systems deliver long-term reporting, but fall short on real-time reporting. Here, too, flow-correlated technology is helpful.

Summary

A distributed network application management solution based on flow-correlated data collection and analysis is an appropriate solution for managing large, complex, distributed networks. It provides real-time coverage at a practical cost, and presents data in its application, user and business context to promote efficient troubleshooting as well as tasks such as planning, baselining, and chargeback. It scales to the largest networks, and allows the efficient division of management responsibilities across several axes: global/regional, server/application/network or C-level/manager. Integration with local, and therefore more easily managed, centralized SNMP systems, can add drill-down to Layer-2 device details for network-specific management.

(1614 words)