# Advanced Network Monitoring Brings Life to the Awareness Plane

*Andreas Kind and Xenofontas Dimitropoulos, IBM Research*
*Spyros Denazis, University of Patras*
*Benoit Claise, Cisco Systems*

## ABSTRACT

The latest advances in traffic measurement, analysis, and modeling play an important role in automatically building and maintaining a distributed intelligent monitoring layer that we describe as the awareness plane. The purpose of this article is to describe the components of this awareness plane in the areas of flexible network measurement, application and relationship discovery, and traffic classification, as well as data aggregation and semantically enriched infrastructure models. We present management services and scenarios, and list the research challenges on the path to integrating the components and making them interoperate for future autonomic service and network management approaches.

## INTRODUCTION

The increasing size, complexity, and dynamic character of network and service infrastructures call for adaptive and self-management functionality as well as increased levels of network manageability [1, 2]. Traditional monitoring systems are no longer able to adequately support administrators with critical tasks in fault, configuration, performance, and security management. As a reaction, the trend is to move away from individual device management toward distributed and autonomous resource control that can be linked to business-level objectives.

Policy-based management [3] is one approach that promises to drive the operation of resources by business-level rules. Rules are expressed using an information model about the computing infrastructure. The dynamic nature of infrastructures in terms of topology, load, and availability requires that the infrastructure model be constantly updated according to the current configuration and operational status. It is also important that the model reflects more than the simple status and configuration of individual devices and components; relationships and dependencies as well as the compound behavior of groups of elements have to be included.

An accurate, semantically enriched model of the network and service infrastructure is not only required for policy-based management, but in general for any attempt to handle the increasing complexity and emerging requirements in future approaches for network and service management based on distributed and autonomic decision-making processes. In this context, the latest advances in traffic monitoring, measurement, and analysis play an important role in automatically building and maintaining an awareness model — a model that captures the infrastructure resources and services, as well as their constraints and dynamic usage. The new methods range from active and passive discovery, traffic metering, and application and service detection to sophisticated aggregation, analysis, and storage of traffic measurement data. The ubiquitous availability of traffic measurement information from routers, switches, servers, and dedicated metering devices has spawned research on analysis techniques that provide the means for populating infrastructure models with much more accurate and semantically enriched information than in the past. Metering at the interface level is combined with end-to-end flow-based measurement and high-speed deep packet inspection. In addition, it is now feasible to configure traffic meters dynamically in order to export traffic information at the exact level needed for a particular monitoring task. Subsequent traffic flow analysis and adaptive storage mechanisms use new aggregation approaches that can accommodate very large streams of measurement data, have means to store relationship information, and can represent behavior modes in multidimensional feature spaces.

Network traffic analysis is crucial for operating networks. It traditionally provides insight into problems and actual usage of links, protocols, servers, applications, and so on. In addition to the benefits it provides at the level of the infrastructure elements, measurement information can be valuable as a basis for analysis on a higher application and work flow level. The reason is that the execution of almost any business application leaves a footprint in network traffic as it involves access to and communication between networked resources. This trend will be further boosted with composite service applications in service-oriented architectures (SOAs)

 0163-6804/08/$25.00 © 2008 IEEE 

and Web 2.0. Figure 1 illustrates how the execution of a business application that involves end-user requests to a Web server and communication between the Web server and other servers (e.g., for authentication) leaves a specific footprint in network traffic. The reoccurring footprint first involves users accessing a Web server, then communication from the Web server to an application server, and finally negotiation with an authentication and a database server. This communication pattern can be identified and characterized by analyzing the sequence of observed flows and their timing and duration characteristics. In this way essential information is obtained for populating, updating, and verifying dependencies of an infrastructure awareness model.

The example in Fig. 1 stems from an enterprise environment. The first priority in enterprise networks is the availability and performance of servers, and the associated services offered to end users. The first priority in traditional provider networks is, however, traffic transit and meeting service level agreements (SLAs) to other customer networks. Despite these differences in network infrastructure and business priorities, the need for a distributed intelligent monitoring layer is shared.

The purpose of this article is to show how state-of-the-art traffic analysis and monitoring techniques can be exploited for modeling complex network and service infrastructures. In that sense, advanced monitoring brings life to an awareness plane that enables a more holistic approach to management based on adaptation and distributed self-management. We identify management scenarios that benefit from the awareness plane (anomaly monitoring, server consolidation, impact analysis) and present the main research challenges in this area, including a critical discussion of the opportunities and limitations of traffic monitoring with the advent of IPv6 and packet encryption.

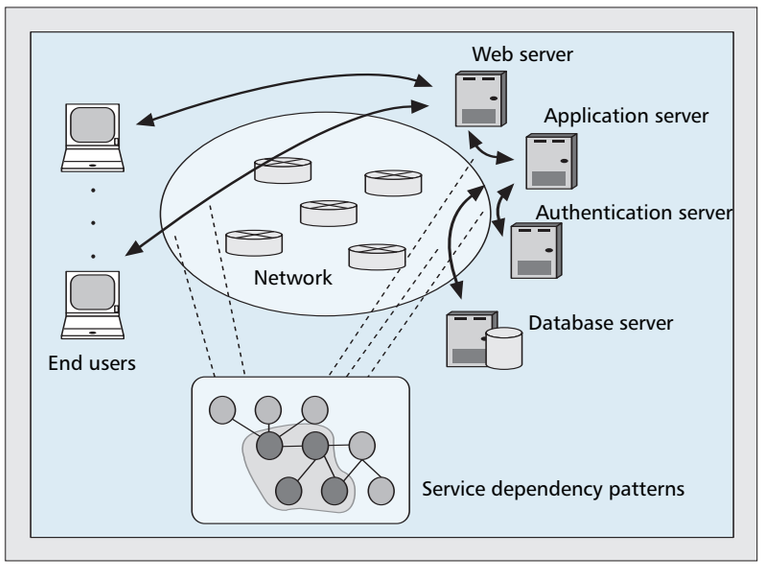


n Figure 1. *Example execution of a business application that leaves a footprint in network traffic.*

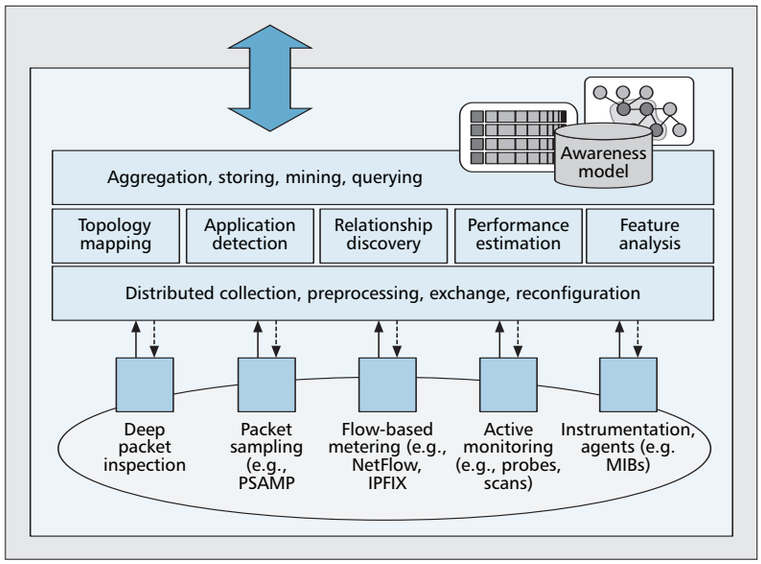


n Figure 2. *The components of the awareness plane.*

## Modeling Network and Service Infrastructures

Finding out the status quo is at the beginning of almost all goal-oriented attempts to manage a process, system, or infrastructure. This also applies to network and service management. With the objective of supporting the business processes, IT managers need to maintain and manage the overall IT infrastructure based on a thorough understanding of the current deployment and operational status. Misconceptions can quickly lead to significant expenses at the business level. An essential method for capturing the deployment and operational status is to create a description or representation of the infrastructure and its status. This description is typically called a model.

Standard methods of building network and service infrastructure models have reached their limits in today's constantly changing environments. In particular, the flexible deployment of multiple virtual resources on a single hardware platform, as is done, for example, with logical operating system partitions, has complicated the modeling task. Also, the increased mobility and heterogeneity of devices, such as laptops and PDAs, creates difficulties in keeping a model up to date.

We advocate the need for evolving existing monitoring approaches toward an awareness plane that consists of distributed traffic measurement and monitoring techniques; dynamic configuration capabilities for traffic data metering, collection, and storage; as well as traffic data analysis for inference and deduction. We use the term *awareness* as we believe that the ability to produce, maintain, request, and discover information is the basis of "being aware." We see the awareness plane as a starting point or subset of the knowledge plane [1] or the idea of autonomous management of communication networks [2]. These research and standardization

*Models are implemented with a combination of relational, object-oriented, and aggregation database approaches. Object-oriented databases simplify the representation of relationship information, and aggregation databases help handling the large sets of monitoring time-series data.*

activities target two complementary scientific areas: semantic modeling and algorithms aimed at information gathering and inference. The combination of these areas enables the awareness plane with the components shown in Fig. 2. Active and passive measurement components, such as deep packet inspection, packet sampling, flow-based metering, server and device instrumentation, and agents as well as infrastructure scans, are used to provide primary traffic and infrastructure information. Clearly, not all types of primary measurement data are typically available in an infrastructure. Primary data is collected and preprocessed (e.g., filtered or pre-aggregated) in a distributed manner. If a reconfiguration is necessary to achieve better results, measurement processes can be adjusted dynamically, and data is exchanged between the processes.

In the next step, various analysis components take the collected measurement data as input. The most important fields for analysis are topology mapping, application detection, relationship discovery, performance estimation, and feature analysis. Extensive amounts of monitoring data generated from large provider and enterprise networks require dedicated aggregation and storing capabilities that, for instance, automatically reduce data fidelity over time, but keep high accuracy where necessary. Data that has already been in the database for several months can be stored with lower resolution than data that is just some minutes old. This approach enables the setting of upper limits on disk and memory consumption even if the database is constantly being filled with new data. Other aggregation methods exploit the locality of particular usage modes by automatically grouping traffic into multidimensional traffic clusters [4].

At the top in Fig. 2, aggregated storage mechanisms are combined with information models, such as the common information model (CIM). Recently, there is a push toward enhancing legacy information models with semantically rich and extensible information. DENng is a typical example of such efforts to raise the level of network manageability and openness. Models are implemented with a combination of relational, object-oriented, and aggregation database approaches. Object-oriented databases simplify the representation of relationship information, and aggregation databases help handling the large sets of monitoring time-series data.

## Measurement and Analysis Components of the Awareness Plane

This section lists the key measurement and analysis components that are instrumental in building and maintaining an awareness infrastructure model.

### Deep Packet Inspection and Sampling

Since traffic is effectively transmitted as packets, the most detailed information can be obtained from packet header capturing and deep packet inspection. An example of a feature that can only be extracted on a packet level is the interpacket arrival time. Deep packet inspection is mostly performed on dedicated hardware appliances connected to a network tap or mirror port of a switch. Today's challenge is to parse header information, including parts of the payload, at wire speeds above 10 Gb/s using regular expressions. Here it is important to use a compact representation of regular expressions and ensure a balanced execution on modern chip technology [5].

With higher link speeds, more interfaces, and generally much more traffic, it is necessary to control resource usage at packet monitoring devices by means of sampling and aggregation. However, in the past, the effects of these preprocessing steps on the resulting traffic models were often not fully understood (e.g., sampling bias toward large flows). Much research has been conducted over the past years to clarify the impact of aggregation and sampling of collected data [6].

Deep packet inspection appliances can be difficult to deploy in a network infrastructure, because to achieve high traffic coverage, devices would be necessary for many links. Deployment may also be prohibited or constrained for privacy reasons.

### Flow-based Metering

In the past, access to management information bases (MIBs) via Simple Network Management Protocol (SNMP) was the only scalable way of collecting data. Today, flow information export based on NetFlow/IPFIX [7] has widely replaced SNMP polling, and provides a detailed end-to-end view that goes far beyond the link- and interface-level information maintained in MIBs. Cisco's Flexible NetFlow is the new flow metering architecture that provides arbitrary key and non-key field combinations. All parts of the IP header and even payload sections can be tracked. It enables the system administrator to tailor the flow record definitions for specific applications. In addition, Flexible NetFlow allows the dynamic creation of new flow monitors to focus on specific network information.

One of the prerequisites for a flexible NetFlow architecture is the NetFlow Services Export Version 9, or the standardized version of it: IPFIX. Both protocols are based on the separate export of templates and flow records. A template is first exported, specifying to the collector which information elements compose the flow records. Once the template has been received, the collector has all the information necessary to decode the subsequent flow records. Flexibility and extensibility are the advantages of template-based flow information export. Indeed, while the export flexibility is offered by the template mechanism, the export extensibility is as easy as adding new information elements to the information model: no changes of the protocol specifications are required.

### Analysis of Large Data Sets

Large provider and enterprise networks generate enormous amounts of monitoring data. For example, few years ago, the network of a large provider was reported to generate 500 Gbytes of NetFlow data per day [8]. The sheer volume of

data has triggered a number of research activities aimed at collecting and analyzing traffic measurements. Efficient algorithms are necessary for computing various statistics. Such algorithms typically treat online traffic measurements as a data stream and process incoming data as it arrives. For example, certain data streaming algorithms compute network traffic heavy hitter (i.e., the largest traffic flows in a network). A classical method for finding heavy hitters is to use a counter to track the size of each flow and sort the flows by their size. However, large network traffic datasets have many distinct flows, which require storing many counters and consume significant memory resources. For this reason, data streaming algorithms for finding heavy hitters do not use a counter for each flow, but instead use a small amount of memory and intelligent techniques to approximate the heavy hitters of a data stream on the fly (e.g., [9]). In this manner it is not necessary to sort a large number of counters, which is an expensive operation, and memory usage is reduced. Data streaming algorithms are currently available for estimating certain statistics, such as the distribution of flow sizes and entropy of traffic data, and are generally required for most types of analysis on large network datasets.

### Relationship Discovery

Direct traffic relationship information, in the sense that a source sends packets to a destination over a specific protocol, is provided with both packet inspection and flow-based metering as well as with some forms of server instrumentation. In many cases direct relationship information already reveals the role of the two end systems involved. An end system is identified as a server when the peer system initiated the communication or known service ports are accessed on the system. Application protocol analysis during deep packet inspection can provide the best evidence that two end systems communicate according to a certain relationship. When an analysis runs for extended time periods, it can also be discovered whether clients use a secondary server for load balancing or failover reasons.

An important question is how servers and services depend on each other and how they support an entire business application, such as a customer relationship management (CRM) or financial accounting system. Such dependency information can be identified by detecting indirect traffic relationships. One way of detecting indirect traffic relationships is to identify flow correlations; that is, flow pairs (or even flow chains) that occur significantly more often than other flow pairs (or flow chains) do [10]. A flow pair is likely to occur more often if the service related to one of the flows in the flow pair requires the service of the second flow in the pair. An example of such a flow pair is a flow representing database access and a flow representing directory access for authentication. The reason for the correlation of these two flows might be that users have to authenticate with the directory server prior to accessing a database. Figure 3 illustrates that the dependency of a database service on an authentication service can be detected from the traffic observation that every flow triggered by an end-user database request is quickly followed by a flow from the database server to the authentication server.

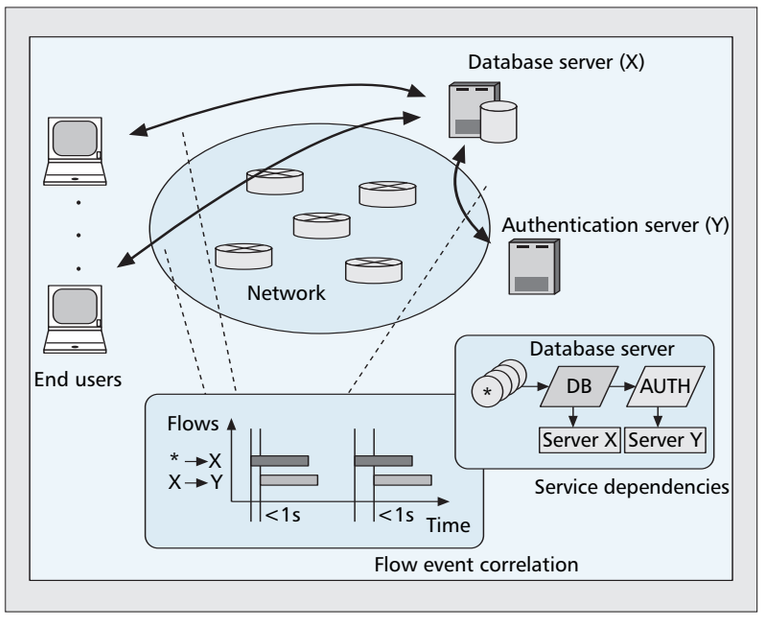


n Figure 3. *Temporal flow event correlation indicates service dependence.*

Practitioners start using auto-discovery tools that combine relationship discovery with discovery mechanisms for hosts, servers, printers, routers, and other devices. The tools apply active and passive techniques, such as fingerprinting of operating systems, routing protocol analysis, and network scanning.

### Network Topology Mapping

Yet another important network management function is to map and monitor the IP-level topology of a network. The topology of interest is in essence a graph in which nodes represent routers or hosts with one or more IP addresses and edges represent links between two interfaces. The classical and most widely used approach to discovering the topology of an enterprise network uses SNMP. A discovery agent sends recursive SNMP queries to the routers of a network to determine the neighboring routers and hosts. The agent starts with the default router of the local host, and uses the collected information to build and maintain a network topology graph. The SNMP-based approach is the simplest and most accurate. However, SNMP is not always enabled, or access may be limited. For this reason other approaches have been investigated for discovering the topology of both local and remote networks. Traceroute-based techniques [11] rely on path probing from a set of measurement points to discover network topology. Traceroute is executed typically to a large set of selected destination IP addresses. The data on the router-level paths collected is then combined to construct a network topology. The main advantage of traceroute is its independence from credentials and agents for discovering network links. It can be used for mapping both local and remote networks. However, it has a number of notable limi-

*The analysis of attacks is a difficult task for network administrators, mainly because the information available during and after an attack does not have the right level of depth to protect against the attack and to determine its root cause.*

tations. First, traceroute relies on Internet Control Message Protocol (ICMP) replies that can be disabled. In addition, IP addresses corresponding to different interfaces of a router need to be identified and merged into a single graph node. Identifying such IP addresses is called alias resolution, and is a very challenging and not yet completely resolved research problem (see, e.g., [12]). Finally, other topology mapping techniques rely on information from routing protocols, such as Open Shortest Path First (OSPF), that build and maintain a complete image of the topology of a network.

### Traffic Classification

Knowing the destination IP addresses and port numbers is not sufficient for identifying the application associated with a traffic flow, as ports can be dynamically chosen, or http tunneling may be used to bypass firewalls. For example, an ongoing challenge is to isolate file sharing and voice over IP traffic flows. Sophisticated traffic analysis techniques are necessary and have been the subject of studies in the past few years. These techniques try to add different semantic annotations to network traffic, such as the application associated with a flow or the benign or malicious nature of certain packets, in order to increase awareness of what is happening on a network and thus extract useful information for different network management tasks. Traffic analysis techniques range from simply inspecting destination port numbers and typical traffic volumes over time to more advanced approaches that may search for packet signatures or model the behavior of hosts, servers, and social user groups [13]. Also, components for network-based application recognition (e.g., NBAR) are added in today's routers.

## Management Services and Scenarios

In this section we provide examples of management services and scenarios that benefit from advanced network monitoring and the resulting awareness models.

### Anomaly Detection

Despite the widespread deployment of firewalls and anomaly/intrusion detection systems, IT infrastructures are still threatened by network attacks. The main reason is that yet unknown types of attacks are difficult to sense with signature-based approaches, whereas behavioral approaches, which observe network-wide feature changes, often suffer from high false positive rates.

The analysis of attacks is a difficult task for network administrators, mainly because the information available during and after an attack does not have the right level of depth to protect against the attack and determine its root cause. In default operation it is not possible to collect and store traffic information at the finest level of detail. For this, all traffic in a network (and in fact also in all connected networks) would have to be stored for some period of time — which is clearly not feasible. Instead, administrators try to reap the traffic information relevant to an attack from available logs and traces. If the attack is still ongoing, a zoom into the particular part of the traffic related to an attack can be started. Typically, such a zoom needs multiple refinements and iterations.

Deep packet inspection and flow-based metering tools are particularly important for performing traffic zooms. Packet capturing provides the highest level of detail, but cannot be applied to large amounts of traffic because of long analysis times and the need to use switch port mirroring or extra hardware equipment for every link to be observed. A typical approach is thus to narrow down the traffic by analyzing flow information and, in a later step, to use packet capturing. With Flexible NetFlow, flow monitoring can be configured to capture a wide range of information starting from aggregate flow-level data down to filtered packet headers and parts of their payloads. In the scenario of an attack, a traffic meter at an observation point can be dynamically extended with a new monitor and specific key fields so that the exported flow information contains the information most relevant to the attack.

Future security systems are likely to offer proactive zooming capabilities, which will promote the availability of new data relevant to attacks. This data will be generated automatically in response to exogenous events, tagged with relevant annotations by analysis and inference modules, and stored using an appropriate information model. After some time, the detailed information about an alleged anomaly can be removed from the model.

The combination of service relationship information and detailed traffic flow information provided by dynamic zooming can be valuable for advanced root cause analysis. If the model stores, for instance, detailed traffic flow information as a consequence of an observed overload situation at a server (e.g., database) and the server is known to depend on other services (e.g., authentication), the load on the server hosting the other service should be checked as it may, in fact, cause performance problems.

### Server Consolidation

Server consolidation is an approach to reduce the number of underutilized physical servers and server locations, with the aim of improving the overall operational efficiency in an IT infrastructure. An important aspect to take into account during the planning of the consolidation process is how servers relate to each other. If a particular server is tightly coupled in functionality with another server, it makes sense to collocate both servers in the same data center or as logical servers in the same virtualized environment.

A great help for a smooth transition process is thus an awareness layer that provides an information model with relationships between servers, the service dependencies, and the variation of workload volumes associated with all dependencies. A simple flow analysis of the direct traffic relationships will typically lead to a highly connected relationship graph with low confidence regarding the correctness of the identified services. Application detection and traffic classifica-

tion techniques are therefore required to discover the relevant traffic relationships that constitute tightly coupled groups of physical and logical resources. These groups are then mapped into a new infrastructure taking constraints into account such as service performance, energy efficiency, and rent of property space.

### Impact Analysis

A typical approach in inventory management is to physically inspect the servers as deployed in a datacenter or server room. Such visits are costly and cannot be done in virtualized environments. Another possibility is to track IT components as financial assets. However, often the financial data are recorded from a project perspective and only once (i.e., at the time of purchase). The current usage of these IT components may be very different from their original usage. The most successful inventory approach is to build IT models from information provided by software agents on the devices. The configuration and performance information is accessed regularly via instrumentation and maintained in a model, for instance, implemented as a configuration management database (CMDB). The aim is to have a single repository that integrates information from problem, change, and asset databases. It contains all technical, ownership, and relationship information necessary for service delivery and support.

A typical use case of an infrastructure model in the form of a CMDB is to enable a better estimate of the impact of changes. The more information available on the effective asset relationships, such as dynamic server dependencies, the lower the business risks associated with any modification of an existing deployment configuration. For example, if a change is to be made to a given server that requires some downtime, the support team will be able to determine all the applications and processes that will be affected during downtime. This will enable better planning of the change and allow the notification of all affected users in advance.

## Research Challenges

More research is required to equip the awareness plane with even better techniques and information models. Several topic areas have already been identified in the network management research communities [14]. With respect to network monitoring and measurements, we regard the following challenges as important.

A problem stemming from the sheer volume of traffic measurements is the significant bandwidth resources consumed for transporting measurement data across the backbone of a network to a central collection point. To avoid obstructing the flow of normal network traffic, network administrators collect traffic measurements in multiple regional collection and storage facilities. For this reason, it is often necessary to use distributed processing algorithms to analyze and cross-correlate traffic measurements that reside in different locations of a network. The goal of such distributed algorithms is to compute desired statistics efficiently, while reducing or minimizing the volume of communication traffic between the different collection facilities. In this context, peer-to-peer technologies can also be useful for the communication and organization of the distributed local collectors.

A second research challenge is the extraction of meaningful and accurate information from traffic measurements. For instance, network administrators want to know how much bandwidth is consumed by different applications, such as file sharing software. Answering this question requires modeling the behavior of applications, mining patterns of the software, and identifying such patterns in traffic data. Inference or data mining algorithms are necessary for a number of similar problems, such as detecting network attacks, mapping server dependencies, and identifying network configuration problems. One important challenge here is that network traffic patterns change with time, especially as new applications become popular. For this reason, improving the accuracy of inference and data mining algorithms is a moving target that requires consistent and lasting research efforts.

Another open research topic is the integration and interoperability of components of the awareness plane. The presented scenarios assume data is exchanged between the components, and components can be controlled for different use cases. However, an architecture for the awareness plane does not exist yet.

The increase of IPv6 traffic and packet encryption will aggravate some of the techniques discussed in this article. The address space of IPv6 will be too large for successful address scans, and protocol analysis with deep packet inspection will not be possible with packet encryption. However, techniques on packet header information, including flow-based metering, will still generate valuable measurement data. In the presence of IPv6 and packet encryption, the general research focus is expected to shift toward tracking behavioral, social, and functional traffic patterns.

*Techniques on packet header information, including flow-based metering, will still generate valuable measurement data. In the presence of IPv6 and packet encryption, the general research focus is expected to shift towards tracking behavioral, social, and functional traffic patterns.*

## Conclusions

In this perspective on research directions in network monitoring, we discuss how recent advances in network monitoring have established an important basis for the management of large, complex, and dynamic network and service infrastructures. Our plea is to move away from the ad hoc methods, tools, devices, and protocols for infrastructure monitoring, and transform monitoring into a key network service by building and maintaining an awareness model of the network and service infrastructure. To justify such an objective, we have discussed some representative traffic monitoring components together with scenarios from enterprise and provider networks.

Future work will have to be dedicated to the integration of the components as well as the existing and emerging monitoring technologies, practices, and methods. The envisioned management scenarios assume data is exchanged between the components, and components are combined dynamically in an ad hoc manner. Middleware approaches based on Web services, peer-to-peer exchange, and publish/subscribe systems may be applied in this context. The

*The resulting new monitoring architecture can be conceived of as an independent and critical component of the future Internet as it will allow the seamless evolution of monitoring whenever new requirements will be introduced.*

resulting new monitoring architecture can be conceived of as an independent and critical component of the future Internet as it will allow the seamless evolution of monitoring whenever new requirements are introduced.

## REFERENCES


[1] D.D. Clark *et al.*, "A Knowledge Plane for the Internet," *Proc. SIGCOMM '03*, Aug. 2003, pp. 3–10.
[2] B. Jennings *et al.*, "Towards Autonomic Management of Communications Networks," *IEEE Commun. Mag.*, vol. 45, no. 10, Oct. 2007, pp. 112–21.
[3] M. S. Sloman, "Policy Driven Management for Distributed Systems," *J. Network and Sys. Mgmt.*, vol. 2, no. 4, 1994, pp. 333–60.
[4] C. Estan, S. Savage, and G. Varghese, "Automatically Inferring Patterns of Resource Consumption in Network Traffic," *Proc. SIGCOMM '03*, Aug. 2003, pp. 137–48.
[5] S. Kumar *et al.*, "Algorithms to Accelerate Multiple Regular Expressions Matching for Deep Packet Inspection," *Proc. SIGCOMM '06 Conf. Apps., Technologies, Architectures, and Protocols for Comp. Commun.*, Pisa, Italy, Sept. 2006, pp. 339–50.
[6] T. Zseby, T. Hirsch, and B. Claise, "Packet Sampling for Flow Accounting: Challenges and Limitations," *Proc. Springer Active and Passive Measurements Conf.*, Cleveland, OH, Apr. 2008.
[7] B. Claise, Ed., "Specification of the IPFIX Protocol for the Exchange of IP Traffic Flow Information," IETF RFC 5101, Jan. 2008.
[8] C. Cranor *et al.*, "Gigascope: A Stream Database for Network Applications," *Proc. SIGMOD Int'l. Conf. Mgmt. of Data*, June 2003, pp. 647–51.
[9] X. Dimitropoulos, P. Hurley, A. Kind, "Probabilistic Lossy Counting: An Efficient Algorithm for Finding Heavy Hitters," *SIGCOMM Comp. Commun. Rev.*, vol 38, no 1, Jan. 2008, pp. 7–16.
[10] A. Kind, D. Gantenbein, and H. Etoh, "Relationship Discovery with NetFlow to Enable Business-Driven IT Management," *Proc. IEEE/IFIP Int'l. Wksp. Business-Driven IT Mgmt.*, Apr. 2006, pp. 63–70.
[11] K. Claffy, T. E. Monk, and D. McRobb, "Internet Tomography," *Nature*, Jan. 1999; http://www.nature.com/nature/webmatters/tomog/tomog.html.
[12] Y. Hyun, A. Broido, and K. Claffy, "On Third-Party Addresses in Traceroute Paths," *Proc. Springer Active and Passive Measurements Wksp.*, La Jolla, CA, Apr. 2003.
[13] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "BLINC: Multilevel Traffic Classification in the Dark," *Proc. ACM SIGCOMM*, Philadelphia, PA, Aug. 2005, pp. 229–40.
[14] A. Pras *et al.*, "Key Research Challenges in Network Management," *IEEE Commun. Mag.*, vol. 45, no. 10, Oct. 2007, pp. 104–10.


## ADDITIONAL READING


[1] P. Yalagandula *et al.*, "S3: A Scalable Sensing Service for Monitoring Large Networked Systems," *Proc. ACM SIGCOMM Wksp. Internet Network Mgmt.*, Pisa, Italy, Sept. 2006, pp. 71–76.


## BIOGRAPHIES

ANDREAS KIND (ank@zurich.ibm.com) joined IBM Research in 2000 and leads the systems management activities at the IBM Zurich Research Laboratory. He received his Ph.D. from the University of Bath, United Kingdom, in 1998, and worked on open programmable networks and network processors in the past. From 1998 to 2000 he was research project manager at NEC Europe. His current scientific focus is on network traffic analysis and visualization.

SPYROS DENAZIS (sdena@ece.upatras.gr) received his B.Sc. degree in mathematics from the Department of Mathematics, University of Ioannina, Greece, in 1987, and in 1993 he acquired his Ph.D. in computer science from the University of Bradford, United Kingdom, working on performance modeling and evaluation of computer networks. In 1996 he joined Intracom S.A., Greece, as project leader. In 1998 he joined Hitachi Europe Ltd., where he worked until 2003. Since 2003 he has been an assistant professor at the Department of Electrical and Computer Engineering, University of Patras, Greece.

XENOFONTAS DIMITROPOULOS (xed@zurich.ibm.com) is a postdoctoral fellow in the Computer Science Department of the IBM Zurich Research Laboratory. He received his Ph.D. and M.Sc. degrees from the Georgia Institute of Technology. His main research interests are interdomain routing, traffic flow measurements, modeling of network topologies, and scalable simulation software. He is a Fulbright Fellow, and has received a Best Paper Award and several other distinctions for his research.

BENOIT CLAISE (bclaise@cisco.com) is a Cisco Distinguished Engineer working as an architect for embedded management and device instrumentation. His area of expertise includes accounting, performance, and fault management. He is a contributor to the NetFlow standardization in the IETF in the IPFIX and PSAMP Working Groups. He joined Cisco Systems in 1996 as a customer support engineer in the Technical Assistance Center network management team, and became an escalation engineer before joining the engineering team. He is the author of the ciscopress book *Network Management: Accounting and Performance Strategies*.