



THE LIPPIS REPORT

---

## Application Intelligence

A New Network Service

By

Nicholas John Lippis III  
Publisher, The Lippis Report

April 2007

## **Abstract:**

---

Networks have traditionally supported applications by providing a transport service. Over the years this transport service was optimized through quality of service mechanisms in an effort to prioritize one set of applications over another. After more than ten years of web development and investment, the type and number of applications have increased dramatically, driving new application support requirements within networks. A new class of network service called Application Intelligence has emerged to address a series of problems and provide IT leaders with tools to optimize application performance while increasing user experience. In this paper we introduce the concept of Application Intelligence as a strategy to solve the networked application performance problems faced by most business and IT leaders.

## Table of Contents

---

|   |    |
|---|----|
| ➤ Executive Summary   | 4  |
| ➤ A Vision for Application Intelligence                                       | 5  |
| ➤ Problems That IT Leaders Are Trying To Solve                                | 6  |
| ➤ Application Intelligence Placement  | 7  |
| ➤ The Intelligent Edge: Application Intelligence Surrounds the Campus Network | 8  |
| ➤ Positive Network Design Implications Associated with the Intelligent Edge   | 9  |
| ➤ The Intelligent DMZ   | 10 |
| ➤ Appliance Overlay or Intelligent Network Approach to AI Deployment          | 10 |
| ➤ Implementing Application Intelligence                                       | 11 |
| ➤ Application Intelligence Solution Selection                                 | 13 |
| ➤ Summary   | 13 |
| ➤ About Nick Lippis   | 15 |

## Executive Summary

---

IT leaders develop applications to support business process. Business applications are significant contributors to business enablement and agility to address markets while creating competitive barriers of entry. Business leaders have commissioned large investments in computer, storage and networking infrastructure to support the automation of their business process through IT applications. To optimize business application performance and security, a new network service called Application Intelligence has emerged.

IT leaders are positioning their networks as a strategic business platform. Application Intelligence is a network service which provides IT leaders with a tool to assist them as the line between applications and networks blurs. Application Intelligence peers deep into data packets to identify and distinguish applications well within HTTP so that they are classified and assigned priorities. This classification process enables IT leaders to police applications in an effort to reduce or eliminate unauthorized applications from consuming corporate assets, reducing employee productivity and forcing non-compliance of various regulations, legislation and Presidential Decision Directives (PDDs). New visibility into application behavior over networks is afforded by Application Intelligence providing IT leaders with a tool to aid in the envisioning, design, implementation and management of applications and networks.

Application Intelligence has become available at the right time, as new application delivery models have emerged, representing unknown and unforeseen traffic patterns and network loads. Structured IT applications are converging upon Web services, which make it difficult to distinguish mission critical from recreational applications. Unstructured applications based upon Enterprise 2.0 technologies like RSS, AJAX, mashups, wikis and SOA network behavior is unknown. Thin client and back-end-based application delivery is yet another network unknown. Add convergence or unified communication and TelePresence real-time network requirements and Application Intelligence provides the application view to manage and navigate through this uncharted industry phase of application delivery.

Application Intelligence is being deployed at the network access and the DMZ, creating an Intelligent Edge and Intelligent DMZ. The Intelligent Edge surrounds a network at the wiring closet level, classifying applications, assigning priorities, defending against day zero exploits and providing IT executives with application and network behavior information. The Intelligent DMZ provides protection and defense for web servers against attackers while alerting operations of suspicious traffic flows. Application Intelligence provides insight into internet ingress and egress traffic flows and policing of application types, eliminating unsupported applications before they traverse over the internet optimizing internet bandwidth for mission critical applications and appropriate use. The Intelligent Edge and DMZ prioritize mission critical applications preserving IT investments and optimizing business process delivering corporate material benefit.

As Application Intelligence's value is greatest when deployed pervasively, a five stage Application Intelligence implementation process is offered. Based upon preliminary empirical data, i.e., customer experience, the cost to equip an average campus network with Application Intelligence is 14% of capital spend associated with a total refresh of network switches. Application Intelligence will be bundled with new switch procurements eliminating the need for separate Application Intelligence budget development. With Application Intelligence spend at approximately 14% of switch acquisition cost, IT leaders gain the benefits of application classification, visibility, policing and end-to-end QoS while business leaders gain increased security, regulatory compliance and a competitive advantage through improved business process performance and optimization.

## A Vision For Application Intelligence

---

IT leaders develop applications to support business process. Business applications are significant contributors to business enablement and agility to address markets while creating competitive barriers of entry. Business leaders have commissioned large investments in computer, storage and networking infrastructure to support the automation of their business process through IT applications.

Networks have traditionally supported applications by providing a transport service. More than ten years ago this transport service was optimized through primitive network ports and IP address-based quality of service (QoS) mechanisms in an effort to prioritize one set of applications over another. Technologies such as traffic shaping or rate limiting were introduced to throttle applications up and down. Protocols such as Differentiated Services (DiffServ), Resource reSerVation Protocol (RSVP), and Type of Service (TOS) are in production within switches and routers to mark packets according to the type of service required. In response to these markings, routers and switches use various queuing strategies to tailor performance to requirements. In data centers application or content switching (load balancing, application classification and acceleration plus network security) and message management devices allow networks to improve application performance.

But while these tools and techniques are powerful, they are not systemic or supported across an entire corporate network, leaving large gaps or blind spots to application performance. IT added new applications as well as bandwidth to accommodate the new load. Adding bandwidth was the knee jerk reaction to application performance woes. This is just not the case in modern day corporate networks.

Networks, by definition, connect all IT resources and have become a strategic business platform. An important platform attribute is broad and deep application recognition and performance support. It is fundamental that a broad portfolio of underlying network technology, which enables increased application performance and security, be resident within the network fabric. An application network service can optimize an application life cycle process from application development, test, roll out, production and maintenance. For example, in data centers, networks increasingly understand content and message structure, enabling them to make decisions and perform network operations such as content-based routing and load balancing, optimizing application performance delivery.

To deliver and optimize application performance, network-based application services need to be end-to-end with the network recognizing applications and providing service at ingress and egress points. The overall network will respond not just to prioritizing applications but delivering a set of underlying services to ensure that application performance is of the highest order. This means that networks will recognize applications, signatures and protocols at the point of entry, classify them, mark them with pre-defined priority which network devices understand and enforce them as the application flows across the network independent of its final location, be it to a data center, an end-point, over the wide area, through the campus or to a branch office. This is application intelligence and it will wrap around an entire network.

After more than ten years of web development and investment, the type and number of applications have increased dramatically, driving new application support requirements within networks. In fact, it's not uncommon that most corporate networks support over 600 applications; some IT provided while many others being consumer-based applications such as Instant Messaging (IM), Skype, BitTorrent, etc. A new class of network service called Application Intelligence has emerged to address a series of application problems and provide IT leaders with tools to optimize application performance while increasing user experience.

## Problems That IT Leaders Are Trying To Solve

---

Most IT leaders are confronted with the management of both supported and unsupported applications flowing across their networks. Application Intelligence is a network service, which addresses a series of problems that degrade application performance. These problems can be classified as the following:

**Application Classification** is the ability to assign quality of service (QoS) to each application identified. Sounds simple enough; however, performing classification at a line rate of a gigabit per second speeds across hundreds of applications for thousands to tens of thousands of end-points is daunting. Traditional application classification was based upon network port plus IP address source and destination pair, which is not granular enough to identify the application. Today's application intelligence dives deep into the data packet to identify applications and assign QoS control. Application Intelligence is performing this task at network access or wide area. For QoS, best practice is to mark traffic at the edge of the network in access closets and carry this QoS marking from end to end across the corporate network.

Converged applications such as Unified Communications, IP video and TelePresence is a large driver for application classification, as these applications require real-time performance. As most application development has shifted to a web environment, the ability for Application Intelligence to perform hardware-based deep packet inspection is the technique enabling the classification of different HTTP traffic versus port 80 or URL-based traffic, for example. This is important on two dimensions. First, it offers the ability to classify a large number of applications. Second, as applications are moved between servers, QoS marking polices and classification maps do not need to change as they do now with port and IP address classification, saving network operations time and task.

**Application Policing** is the ability to allow or restrict applications from flowing across the network business platform. With the huge influx of recreational or consumer-oriented applications such as BitTorrent, Kazaa, Skype, Gnutella, Grokster, various IM services, etc., the ability to police the use of these applications on corporate networks is a requirement. This is due to some recreational application network behavior, causing an organization to be out of compliance with regulations, consume large percentages of bandwidth, and/or reduce end-user productivity. For example, Skype, Kazaa, Grokster and other peer-to-peer (P2P) applications create a SuperNode where a computer becomes a relay node within a corporate network, passing traffic on behalf of users with no corporate affiliation of their own, consuming resources and creating security vulnerabilities. Application policing is engaged after classification and QoS marking is complete. Non-mission critical applications may be dropped or their QoS priority set low.

**Miscreant Application Mitigation** is the identification of exploit signatures at network ingress and egress so that network and security operations may respond to day zero attack vulnerabilities. Application policing is a security measure which Application Intelligence affords. But Application Intelligence goes beyond policing to signature identification of exploits such as viruses, worms, Denial of Service (DoS) attacks, etc. Application Intelligence placed close to users in access switches provides the best defense to stop exploits from propagating throughout the network and infecting corporate assets. Application Intelligence placed at data center boundaries and in the DMZ protects servers against attackers and provides insight into corporate DMZ traffic flows. Operational staff would be alerted during the protocol discovery phase of application/traffic signatures not included in DMZ security policy, prompting investigation. In short, Application Intelligence analyzes application signatures and if a miscreant application is identified Application Intelligence either throttles the application down, drops the application and/or keeps a log of the action for audit purposes. Note that the network security advantages of Application Intelligence are often its economic and benefit justification.

**Application Behavior Visibility** is the ability to view and gain insight into how applications behave while flowing through a corporate network. How application use impacts bandwidth, server load, internet connections and user experience are data provided by Application Intelligence. Unified Communications injects **real-time** voice, audio, IM and voice messaging while TelePresence adds conferencing to the collaboration experience, all of which are optimized thanks to Application Intelligence's visibility into network behavior. Application behavior visibility is extremely helpful to IT developers, network architects and security operations, providing these groups with a view into application interaction with end-points, servers, storage and networks. This visibility adds value to the entire application life cycle from application development, test, roll out, production and maintenance. Increased visibility into application behavior leads to improved design, user experience and business process.

Apart from traditional structured application development, IT departments are embracing Enterprise 2.0 concepts to deploy free-form social software to increase collaboration through the use of RSS, AJAX, mashups, wikis and SOA. Enterprise 2.0 enables all users to create content they see as relevant while the network makes sharing this content and collaborating possible. Enterprise 2.0 projects represent unknown and unforeseen traffic patterns and network load. The application behavior visibility of Application Intelligence will aid IT departments as they deploy unstructured corporate collaborative applications by providing a view into the behavior of these applications, which IT will use to optimize their deployment.

## Application Intelligence Placement

---

As Application Intelligence is a new network service its placement within a network needs to be explored. One example of Application Intelligence can be found in routers today via Cisco's **Network-Based Application Recognition (NBAR)**. To expand NBAR's functionality reach throughout a corporate network and increase performance to the multi-gigabit level, Cisco has introduced its **Programmable Intelligent Services Accelerator (PISA)**, into the Supervisor Engine 32 for the Catalyst 6500 series of switches. There are appliance-based Application Intelligence solutions as well, from various vendors such as Packeteer, F5, BlueCoat, et al.

With Application Intelligence embedded into the network fabric there are two primary placement options that deliver the greatest value. Those are in access switches closest to users. This option is called the **Intelligent Edge**. The second placement option is in the DMZ called the **Intelligent DMZ**. Note that these options are not mutually exclusive but in fact overlap to maximize Application Intelligence coverage and benefit.

## The Intelligent Edge: Application Intelligence Surrounds the Campus Network

Figure 1 illustrates the Intelligent Edge option. In a typical three-tiered building architecture, wiring closets house access switches, which define the edge of the network while these switches aggregate into a distribution layer of switches. Distribution switches are connected via core switch/routers. Placing Application Intelligence at the edge of the network addresses all of the above mentioned problem statements.




Figure 1: The Intelligent Edge Classifies Applications and Filters Application

**Application Classification:** Classifying applications and marking QoS at the network edge is a best practice to improve application performance as QoS marking can be carried from end to end across the corporate network, from edge to edge.

**Application Policing:** Policing applications at the network edge before they enter the network and propagate is a best practice to limit recreational or consumer-oriented applications such as BitTorrent, Kazaa, Skype, Gnutella, Grokster, IM, etc. By policing applications at the edge a corporation is assured of high performance for mission critical applications and business process.

**Miscreant Application Mitigation:** Identifying exploit signatures for notable worms/viruses and their variations at the network edge delivers day zero defense from malware propagating throughout a network causing damage to IT resources and avoiding the operational spend associated with post-exploit clean up.

**Application Behavior Visibility:** Implementing Application Intelligence from edge to edge provides IT leaders with the greatest visibility into application behavior as observed between end-points.

## **Positive Network Design Implications Associated with the Intelligent Edge**

The Intelligent Edge filters and cleanses traffic flows before they enter the distribution and network core providing a defensive barrier around the network, which allocates network resources for mission critical business applications and process. Further, with a pervasive Intelligent Edge classifying and assigning QoS to applications, wide area network devices need only enforce QoS rather than assign its values. This could eliminate the need for traffic shaping at the wide area and other QoS assignment devices.

The Intelligent Edge will impact data center design over time too, by reducing classification requirements in the data center and even providing load balancing at the edge of the network. Application Intelligence does not remove the requirement for content switching in the data center; however, it may have a tendency to offload some of this functionality as it's widely distributed across the edge of the network.

The Intelligent Edge over time will reposition certain features and functions within a corporate network. For example, classification consumes large compute cycles of existing equipment such as firewalls. As classification of applications is conducted at the edge, every packet is inspected without compromising performance. Application Intelligence can then create a linkage of these inspected packets with a firewall that informs it that these packets have been classified. This linkage enables firewall- Application Intelligence interaction, allowing the firewall or Application Intelligence to take over flow control.

Application Intelligence and firewalls can work together to deliver increased security and application performance. For example, the Intelligent Edge could identify applications that are port hopping, e.g. Skype, and provide this information to the firewall, which may not be capable of performing this identification. The firewall, armed with this information will now be able to apply policy that could block such applications. Hence Application Intelligence at the edge messaging to policy enforcement at the firewall provides a comprehensive application security solution.

With the Intelligent Edge conducting classification and defending against exploits the firewall and IPS can operate at tremendous speeds, freeing the firewall and IPS engines from classifying and filtering signatures for some 20Gbps worth of traffic. 20Gbps is based on a typical campus network made up of ten access layer switches connected at 1Gbs each for up and down links. By distributing classification and prevention to Application Intelligence the firewall and IPS are offloaded as a large part of their computer engine processing is now performed at the edge. The firewall is freed to apply policy, its primary function while the IPS can be focused on protecting the enterprise from more sophisticated exploits.

## The Intelligent DMZ

Figure 2 illustrates the placement of Application Intelligence in the DMZ, which provides protection and defense for web servers against attackers while alerting operations of suspicious traffic flows. The protocol discovery attribute contained in Application Intelligence alerts security and/or network operations to malicious applications, which could prompt investigation. Application Intelligence provides insight into internet ingress and egress traffic flows and policing of application types and URLs allowed on corporate networks. Eliminating or dropping unsupported applications before they traverse over the internet optimizes internet bandwidth for mission critical applications and appropriate use. Stopping exploits flowing into a corporation from the internet through their DMZ mitigates this threat. Traffic classification and shaping of internet traffic at ultra-high speed, in the multiple gigabit per second range, is suitable for internet access links up the OC-48 level.

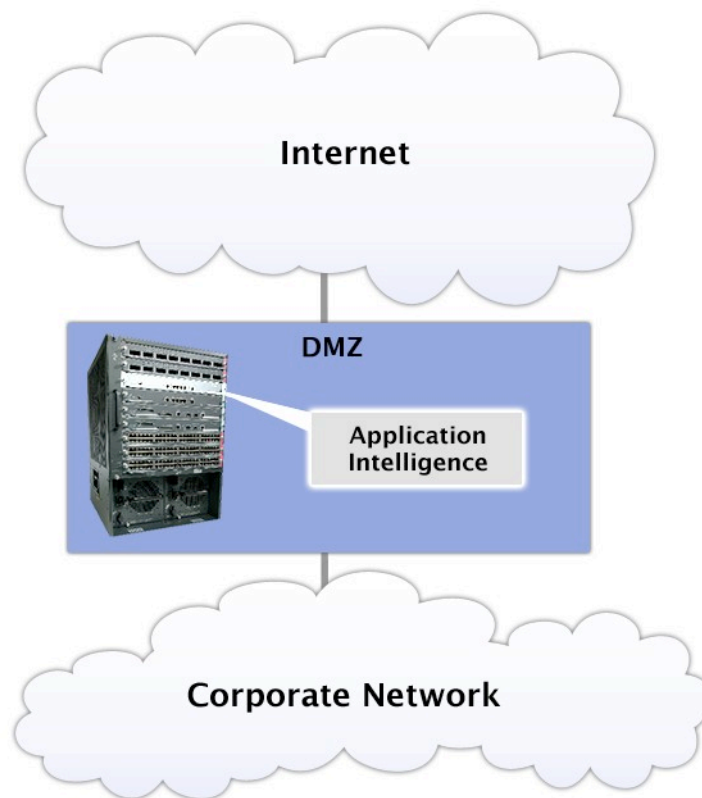


Figure 2: The Intelligent DMZ Filters Application Between Corporation and Internet

## Appliance Overlay or Intelligent Network Approach to AI Deployment

As mentioned above there are vendors offering appliances that offer aspects of Application Intelligence. There is conventional wisdom in the industry that points to appliance fatigue. With every new wave of network technology appliances are introduced with an overlay implementation strategy. There are overlays for wireless LANs, NAC appliances, WAN acceleration, content switching and many more. Each overlay is a separate network which needs to be configured and managed, driving up operational cost. Overlays offer little to no control point interfacing with network infrastructure devices, thereby eliminating the opportunity for collaboration to address a holistic Application Intelligence solution. Further, the number of appliances to address the Intelligent Edge would be daunting. Consider a typical campus

network of ten wiring closet switches, each aggregating 200 users supporting a total of 2000 users with 2x1Gbps uplinks. To bring two gigabits worth of application classification QoS networking plus embedded security for virus mitigation on each of these wiring closet switches would require some twenty plus appliances to configure and manage.

It is for the reasons above that appliances for Application Intelligence be used sparingly. There are network points where appliances could have a positive effect. These points are isolated to data center and DMZ. However, appliances and embedded Application Intelligence from different vendors will not interoperate, message or share data. Application Intelligence as an embedded service is the most advantageous approach as it can blanket the edge of the network and be equipped in data center and DMZ switches offering a holistic Application Intelligence approach with a single control point. The intelligent network approach to Application Intelligence leverages operational staff skills by access Application Intelligence configuration and control points from known management interfaces.

## Implementing Application Intelligence

---

As with any enterprise wide project, implementing Application Intelligence requires planning, preparation, design, implementation, operation and optimization. While implementing both the Intelligent Edge and DMZ require planning, the Intelligent Edge demands more thoughtfulness as its deployment is pervasive. Staging is a best practice in the design and implementation phases of Application Intelligence. Figure 3 illustrates Application Intelligence deployment stages while the text below describes each stage.

**Stage One - Pilot and Test:** The first stage of Application Intelligence deployment is to pilot Application Intelligence in a lab or off-network environment. If corporate security policy concerning application type and use has not been agreed upon with business unit leaders, now is the time for such deliberations and planning. The goals of this stage are to test performance, policy cause and effect and to educate operational staff as to configuration and management. Once staff has gained operational confidence and Application Intelligence functions and features have been thoroughly tested and customized to corporate IT policy Stage Two can commence. Stage One is a one to three month process depending on the size of the organization and deployment scenario of Intelligent Edge, DMZ or both.

**Stage Two - Limited Deployment:** In Stage Two, IT leaders work with business unit managers to identify appropriate limited deployment. IT leaders may consider groups, which are more accommodating to test scenarios. Also it's recommended that consideration be placed upon users who would gain the most value from Application Intelligence. Administrative workers who rely mostly on e-mail and web traffic may not be as demanding as engineering or research and development power users. In short demanding users would be best to test Application Intelligence in limited production as operational staff would gain the most insight into Application Intelligence features.

**Stage Three - Observer and Optimize:** In Stage Three it is likely that operational experience observed with Application Intelligence in production will be different than Stage One. This is due to the discovery nature of Application Intelligence and the fact that operational staff will have a window into their network, which they have not previously had. Operational staff may very well observe application use they did not anticipate or realize the network was supporting. This will require adjustments to Application Intelligence configuration and potentially multiple cycles of optimization so that lab expectations set during Stage One converge during Stage Three. This optimization cycle should be relatively short and measured by days to weeks. It is normal to engage in multiple cycles of optimization and train operational staff in how to respond to new applications, both unsupported and supported, which Application Intelligence detects.

**Stage Four - Event Based Deployment:** Stage Four is an event-based deployment scenario rather than a full infrastructure upgrade. Event-based deployment leverages opportunities such as the opening of a new building(s), merger or acquisition, moves and changes or the retirement of older switches to deploy Application Intelligence. Event-based deployments take longer to achieve than ubiquitous Application Intelligence deployment, but with every addition of Application Intelligence service added to the network infrastructure IT leaders are able to improve user experience and business process.

**Stage Five - Full Deployment:** Stage Five is a pervasive Application Intelligence deployment including the implementation of an Intelligent Edge and DMZ. Based upon preliminary empirical data, i.e., customer experience, the cost to equip an average campus network with Application Intelligence is 14% of capital spend associated with a total refresh of network switches. With Application Intelligence representing approximately 14% of switch acquisition cost, many IT and business leaders will choose to fully deploy Application Intelligence during network infrastructure refresh cycles as well as new switch additions. This is the deployment scenario I recommend, as Application Intelligence cost/benefit value proposition is attractive. With Application Intelligence spend at approximately 14% of switch acquisition cost, IT leaders gain the benefits of application classification, visibility, policing and end-to-end QoS while business leaders gain a competitive advantage through improved business process performance and optimization. Remember acquisition cost is depreciated over some number of years while Application Intelligence benefits are gained immediately after deployment and are permanent. Pervasive Application Intelligence deployment will provide operational tools to increase application performance enterprise wide versus a piece meal approach through appliances or partial deployment.




Figure 3: Five-stage Application Intelligence Deployment Process

## Application Intelligence Solution Selection

---

The following provides key attributes of an Application Intelligence solution. IT leaders should consider the following requirements when designing and acquiring an Application Intelligence solution:

- Deep packet inspection that classifies traffic into the payload especially to distinguish between HTTP traffic
- Two to two and a half Gbps application classification performance so that Application Intelligence can inspect all packets at line rate without compromising application performance. Two to two and a half Gbps is important, as most campus networks are equipped with 1G up and down links
- Integrated or embedded Application Intelligence into network switches and not appliances for the Intelligent Edge
- Integrated Application Intelligence configuration and management into existing management systems. Management consistency is important for operation efficiency, which is usually gained through an embedded Application Intelligence service and lost with appliance-based Application Intelligence deployments
- Virtualized Application Intelligence management and configuration interfaces so that IT developers, Security Operations, Network Operations, etc., are supplied with their own view into Application Intelligence which addresses their respective responsibilities
- A company that can provide support in your location
- A scheme to keep Application Intelligence up to date with current virus signatures so that Application Intelligence can pull down signature files from a central repository

## Summary

---

Application Intelligence is addressing some of the most difficult requirements confronting IT leaders. As corporate networking has become a strategic business platform, business and IT leaders require an understanding of:

- Who is accessing the network
- How are they authenticated
- What type of work do they do
- What is the compliance policy of the application accessing the network
- How can I monitor what applications are running in my network
- How do these attributes relate to the services users are allowed to access in the network
- How do they transverse through the network
- How do I gain a better sense of what applications are emanating from my desktop and how they are utilizing my network
- Where are the resources going
- How do we gain total control, manage, or make sure applications get the right capabilities in the network and preserve user experiences, while the wrong applications can be throttled down or out

Application Intelligence addresses the above concerns and requirements. But the above list is what dominates the minds of most IT leaders today. Over the next few years the requirements will be more demanding. For example, consider that an important meeting is taking place in a TelePresence conference room, which needs to be broadcast to all desktops or pertinent executives and employees. TelePresence is broadcast via IPTV to these authorized desktops. While the IPTV broadcast is taking place, in a typical campus network there are hundreds of applications flowing, supporting thousands to

tens of thousands of users. How does a network operation ensure that the network prioritizes all these applications so that voice services, TelePresence and all real-time applications obtain their share of network resources and priorities? Network resources need to be provisioned all the way to the QoS, bandwidth and service levels to ensure not only that applications are delivered in high performance but that real-time services such as TelePresence are delivered synchronously with voice and video delivered relatively at the same time to preserve this experience.

Previously, IT executives would throw bandwidth and core switch and routing processing at the problem of poor application performance. Others would add bandwidth plus spend large amounts of time implementing QoS mechanisms in both the LAN and WAN. Application Intelligence offers a different approach to campus network design by automating application performance management. Where there are multiple different applications flowing across the network, they can all be logical, meaning that a network is made up of "n" logical networks with Application Intelligence managing these logical entities. Application Intelligence is managing the connection between applications, be they a strong or weak connection, or no connection at all. Application Intelligence allows the application to view the network as though it is matching the network's capabilities with its unique requirements. In short, Application Intelligence allows every application to obtain its fair share of resources, bandwidth, QoS, and Service Level Agreement (SLA) in the presence of all other applications.

## About Nick Lippis

---



Nicholas J. Lippis III is a world-renowned authority on advanced IP networks, communications and their benefits to business objectives. He is currently working with clients developing converged network architecture, which includes IP telephony, secure networks, wireless LANs, internet data centers and storage area networking. He is the chairman and host of the Enterprise IP Communications 2005 Symposium, a conference where corporate network architects and designers learn and share industry best practices. Mr. Lippis hosted thirty-seven sponsors and four hundred and sixty delegates during his Enterprise Networks 2004 conference in Boston.

He has advised numerous Global 2000 firms on network architecture, design, implementation, vendor selection and budgeting, with clients including Barclays Bank, Microsoft, Kaiser Permanente, Sprint, Worldcom, Cigital, Cisco Systems, Nortel Networks, Lucent Technologies, 3Com, Avaya, Eastman Kodak Company, Federal Deposit Insurance Corporation (FDIC), Hughes Aerospace, Liberty Mutual, Schering-Plough and many others. He works exclusively with CIOs and their direct reports. Mr. Lippis possesses a unique perspective of market forces and trends occurring within the computer networking industry derived from his experience with both supply and demand side clients.

Mr. Lippis founded Strategic Networks Consulting, Inc., a well-respected and influential computer networking industry-consulting concern, which was purchased by Softbank/Ziff-Davis in 1996. Mr. Lippis was named one of the top 40 most powerful and influential people in the networking industry by Network World. For nine years Mr. Lippis reached over 120,000 purchasers of networking equipment and services through his monthly column "Lippis on Internetworking" published in Data Communications magazine. He was a contributing editor and columnist for Tele.Com magazine reaching over 80,000 service provider professionals monthly. He currently writes the "Lippis on IP Communications" column for Network World reaching 180,000 in print and 850,000 online. He publishes The Lippis Report, which is distributed to over 360,000 senior IT executives around the world. Mr. Lippis' reach exceeds 1,400,000 readers. He is a frequent keynote speaker at industry events and is widely quoted in the business and industry press.

Mr. Lippis received his Bachelor of Science in Electrical Engineering and his Master of Science in Systems Engineering from Boston University. His Masters' thesis work included selected technical courses and advisors from Massachusetts Institute of Technology on optical communications and computing.