

The Internet Protocol Journal

June 2009

Volume 12, Number 2

A Quarterly Technical Publication for
Internet and Intranet Professionals

In This Issue

From the Editor	1
DNS Caching.....	2
IEEE 802.21	7
Book Review.....	28
Fragments	30
Call for Papers.....	31

FROM THE EDITOR

After many years of using DSL as my only Internet access option from home, I recently upgraded to a broadband solution provided by a cable modem. As a result, I faced the task of renumbering (and partially rewiring) my home network. As you might have guessed, the addressing scheme provided by my new ISP offers *Network Address Translation* (NAT), as well as a small number (5) of fixed IPv4 addresses, the latter at an extra cost as you might expect. I probably should have tried to enable IPv6 just as an experiment, but this task will have to wait for another day. In the meantime, I was pleased to find a relatively user-friendly web interface to the cable modem that allows me to configure numerous parameters, including the range of the *Dynamic Host Configuration Protocol* (DHCP) pool so that certain devices (printers and wireless access points in particular) can have fixed IP addresses for ease of use and configuration. The entire exercise, which took a couple of hours on my very small network, reminded me of what network managers face every day, particularly as they consider the inevitable migration to IPv6. Let me take this opportunity to invite you to share your network management and operations experience, plans for IPv6 migration, and so on. You can send us Letters to the Editor or article proposals. The address, as always, is ipj@cisco.com

The *Domain Name System* (DNS) has been the target of attacks over its many years of existence. In recent years, new attacks have emerged that exploit some of the attributes of the DNS protocol and its implementation. One of the corrective measures is to improve the security of DNS caches. There are several ways to improve cache security, most of which involve changing the protocol. Another way, without changing the protocol, is to reduce the attack surface of your cache by shrinking the number of users of any given cache. Our first article, by Bill Manning, explores this view in more detail.

This journal has covered numerous current and emerging *wireless* technologies such as Bluetooth, Wi-Fi, WiMAX, and mobile cellular systems. In this issue, Esa Piri and Kostas Pentikousis describe *Media-Independent Handovers* (MIH), which allow mobile devices to use different wireless and wired network infrastructures transparently. The protocols associated with operation across such diverse access networks are being standardized by the IEEE 802.21 working group.

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

You can download IPJ
back issues and find
subscription information at:
www.cisco.com/ipj

ISSN 1944-1134

Intermediate DNS Caching as an Attack Vector

by Bill Manning

The *Domain Name System* (DNS) specification calls for the use of *caching*. Caching is expected to improve the overall responsiveness of the system by ensuring that answers to questions are known and stored locally and that the query load placed on the authoritative servers is minimized. Certain presumptions are associated with caches that may no longer hold. This article looks at some of these presumptions and explores some of the problems that emerge when they are violated. Based on our observations, we offer some recommendations on DNS cache best practices and show our results of testing these practices.

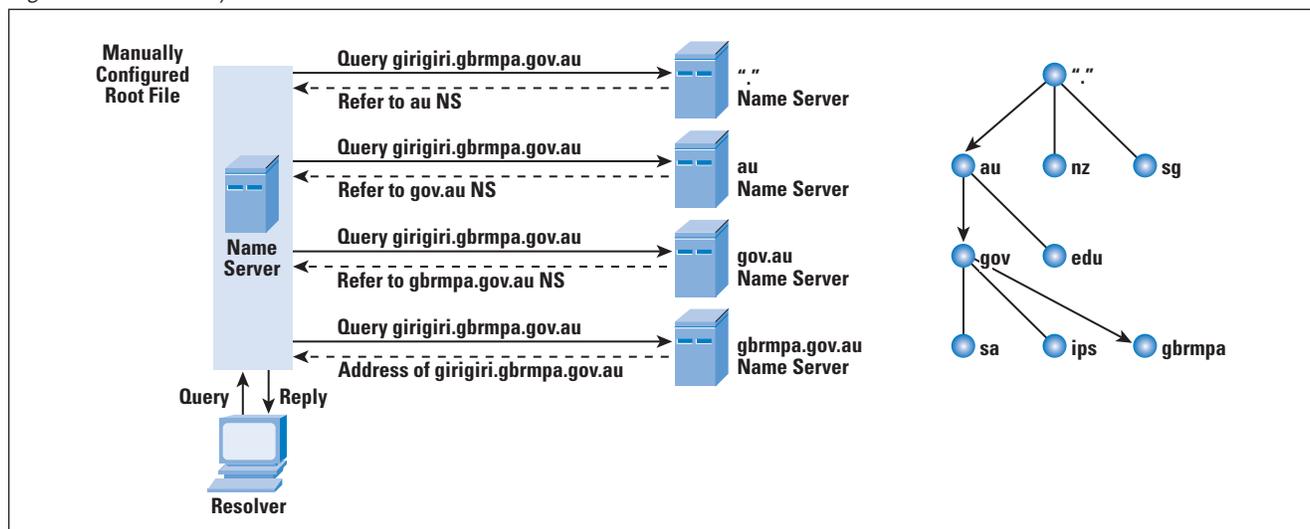
The Problem

A DNS resolver can no longer trust the data it gets—because the data generally comes from nonauthoritative nodes or caches operated by third parties, most of whom have no vested interest in providing accurate data. Removing or bypassing caching from the DNS and going directly to the authoritative servers is considered a fatal flaw because authoritative servers are presumed to have neither the bandwidth nor the processing power to accommodate the perceived demand from a cacheless service. This article looks at the bandwidth and processing capabilities of modern authoritative servers to ascertain the viability of these presumptions. We start by looking briefly at the DNS.

The DNS

The DNS namespace is made visible and useful by nodes publishing authoritative information about the namespace and *resolvers* that send queries about the namespace to these servers. As an optimization, other nodes may act as intermediates or proxies for the authoritative servers for one to many resolvers. These intermediate nodes are called *caching nameservers* or *iterative mode resolvers*. This flow is shown in Figure 1.

Figure 1: DNS Query Flow



Several assumptions about the use and placement of caches have been questioned recently. The simplest is one of placement. A cache works best when the *Round-Trip Time* (RTT) between the resolver and the cache is low. Historically, a cache was placed at traffic aggregation points such as an *Internet Service Provider* (ISP) operating a cache for its clients. With increased mobility of nodes, this presumption is no longer as firm. There are reported cases where resolvers continue to use caches 300 ms away, while an authoritative server is 15 ms away. So if the intent is to reduce network bandwidth, then a cache presuming its client resolvers are all “local” might be misconstrued.

Fixing a resolver to a specific cache does have the benefit of being tied to a known business relationship; for example, using your ISP’s caching service. In contrast, mobile nodes often get an IP address from a provider’s *Dynamic Host Configuration Protocol* (DHCP) servers, which also hand out more “local” caching servers to be used by the mobile node.

This scenario would be fine—as long as the DNS namespace was in fact a coherent, single space. Unfortunately it is not. So-called *Walled-Garden* networks that have their own versions of DNS namespace have been and remain common. In the Internet, there are more and more alternate root hierarchies that diverge from what most think of as “the” root namespace in either subtle or wildly divergent ways. To date, there is no deployed way for a resolver to determine the origin of the data stored in a cache. A resolver then has no way other than verification of the data to know that the locally assigned cache is in fact using the namespace desired. This situation represents one important reason for going back to a well-known cache, even if it is topologically remote. But this assumption may no longer be valid.

ISPs and even some caching service providers are starting to manipulate caches as a means to monetize their operations.^[1] Numerous techniques are in use, from the nominally benign method of using wildcards to more insidious capture and rewrite of NXDOMAIN replies, to outright intentional cache pollution.

In this climate, a resolver should choose its cache carefully. We argue that it is reasonable, in many of today’s environments, to place the cache within 1 ms of the resolver; for example, run a cache on the local node. This argument is an extension of the assertion^[2] that claims that caches are effective for client populations that are about 10 or fewer.

This technique has the added advantage of reducing the “attack surface” by reducing the effect of cache poisoning or rewriting replies to a small handful of nodes. The perceived disadvantage is the increased load on network bandwidth and query load on authoritative servers as the number of caches increases.

The Experiment

Our experiment has two parts: first we looked at authoritative server processing capabilities and then at the bandwidth effects of a larger number of caches.

Authoritative service is generally run on systems with modern software, supporting threading or precomputed responses. Independent testing shows that these stock software solutions can, on current hardware, support query rates in the hundreds of thousands of queries per second.^[3]

A brief survey of authoritative server operators indicates that normal query rates range from 12,000 to 64,000 queries per second.^[4,5,6]

On the surface, this result would indicate that there is enough overhead to be able to process more queries, regardless of how they are originated. Regarding bandwidth, a survey of *Top-Level Domain* (TLD) operators has shown that 92 percent of the delegations have two or more authoritative servers for that data on networks with a minimum uplink bandwidth of 100 Mbps. Selected path characterization from clients to target authoritative servers seems to support our presumption that bandwidth is not of concern.

The DNS was designed to function as a roughly symmetrical transfer of information: a request or query is sent and the reply reflects the query and supplies the answer and additional data. Historically, the request and reply were within the same order of magnitude. Into the future, this model may no longer be valid. With *Domain Name System Security Extensions* (DNSSEC), *IP Version 6* (IPv6), and *Naming Authority Pointer* (NAPTR) records being possible candidates in the *Resource Record set* (RRset), the traffic profile more resembles an HTTP request/response, with a significant amount of data being returned from a simple question.^[7]

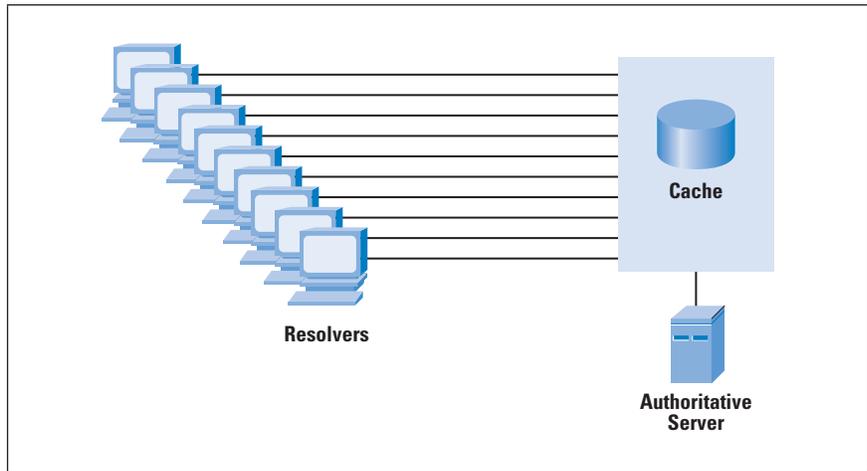
With this information, we can project a worse case in today's environment where a query/reply is about 260 bytes to a worst case in a future environment where a query/reply is about 9 KB, clearly indicating that the amount of bandwidth to authoritative servers needs to grow as new DNS capabilities are deployed, but for the nonce, most have a bandwidth overhead sufficient to absorb a modest change in the number of queries presented.

Modification of the Number of Caching Servers

We began with a cache that serviced 140 stub resolvers on the *University of Southern California's Information Sciences Institute* (USC/ISI) campus in a "normal" dense cache mode (Figure 2).

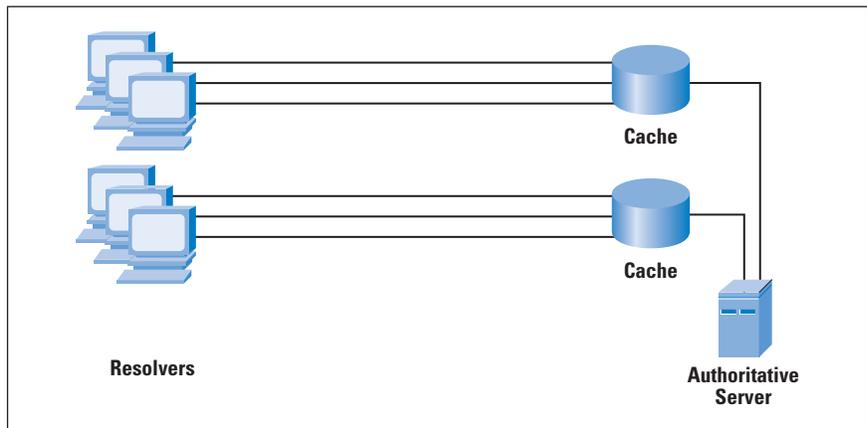
Traffic traces show a distribution of priming queries to 534 authoritative servers in the first 15 minutes of clearing the cache.

Figure 2: Dense Cache



We then added 9 new caches and redistributed the 140 stub resolvers among the 10 caches into a sparse cache mode (Figure 3) and restarted all the caches. In the first 15 minutes, the number of priming queries from each of the caches averaged 61, with a total of 622 unique priming queries for all caches. The number of “duplicate” queries between caches averaged 45. Although the number of queries to the authoritative servers was slightly higher, the results seem to indicate that there is a small but significant difference in each of the caches^[8].

Figure 3: Sparse Cache



Conclusions

Reducing the size of the user population for each cache reduces the attack surface for the DNS overall because we have effectively compartmentalized the threat to a small number of nodes. Generally, restarting a cache for a small number of nodes is considered acceptable, whereas restarting a cache for 10,000 or 100,000 nodes would significantly affect operations.

Moving the cache closer to the resolver increases overall response time and may support better mobility of the node. If validation is also placed with the cache, it is possible to increase the confidence of validation because that information may not have to use DNS protocols to send validation data over untrusted, open networks.

The concept of supporting larger numbers of full DNS servers on more nodes raises concerns, but most systems these days have enough processing power and bandwidth to support this application. Administrative and management processes can be fully automated. Overall, this design complements other, protocol-based attempts to increase DNS integrity.

References

- [1] “Preliminary Report on DNS Response Modification,” 20 June 2008, <http://www.icann.org/en/committees/security/sac032.pdf>
- [2] “DNS Performance and the Effectiveness of Caching,” Jaeyeon Jung, Emil Sit, Hari Balakrishnan, and Robert Morris, *IEEE/ACM Transactions on Networking*, Volume 10, No. 5, pp. 589–603, October 2002.
- [3] <https://www.dns-oarc.net/files/workshop-2006/Dickinson-Performance.pdf>
- [4] “An analysis of Wide-Area Name Server Traffic: A Study of the Internet Domain Name System,” Peter B. Danzig, Katia Obraczka, and Anant Kumar, *ACM SIGCOMM Computer Communications Review*, Volume 22, No. 4, pp. 281–292, 1992.
- [5] “An Analysis of the Queries from Caching Servers to Root Servers, Tsuyoshi Toyono, NTT Laboratories, 2007 OARC Workshop, <https://www.dns-oarc.net/files/dnsops-2007/Toyono-Caching-analysis.pdf>
- [6] RootServer supplied statistics:
<http://h.root-servers.org/>
<http://k.root-servers.org/index.html#stats>
<http://m.root-servers.org/>
- [7] http://snad.ncsl.nist.gov/dnssec/mem_usage.html
<http://snad.ncsl.nist.gov/dnssec/bandwidth.html>
- [8] “Sharp Transition Towards Shared Vocabularies in Multi-agent Systems,” Andrea Baronchelli, Maddalena Felici, Vittorio Loreto, Emanuele Caglioti, and Luc Steels, *Journal of Statistical Mechanics: Theory and Experiment*, 2006, P06014.
<http://www.iop.org/EJ/abstract/1742-5468/2006/06/P06014>

BILL MANNING has been in the network field since 1979, currently with the Keio University, Shonan Fujisawa Campus, and USC/ISI. He has been an IETF Working Group chair and RFC author, and he currently serves on numerous ICANN committees. He is part of the team that runs one of the Internet Root nameservers. E-mail: bmanning@sfc.wide.ad.jp

IEEE 802.21: Media-Independent Handover Services

by Esa Piri and Kostas Pentikousis, VTT Technical Research Centre of Finland

Popular mobile devices now ship with several integrated wired and wireless network interfaces. *Personal Digital Assistants* (PDAs) and smartphones, for example, are increasingly supporting communications through both cellular technologies and *Wireless LANs* (WLANs); laptops typically come with built-in Ethernet, Wi-Fi, and Bluetooth^[1]. As multiaccess devices proliferate, we move closer to a network environment that is often referred to as “*beyond 3G*” (B3G). Key success factors for cellular *third-generation* (3G) communications include better cell capacities, increased data rates, transparent mobility within large geographical areas, and global reachability. For B3G, the next frontier lies beyond transparent mobile connections within the same access technology because users will expect to be globally reachable anytime, anywhere, and remain “*always best-connected*” (ABC)^[2]. In order to select the best possible connectivity option (anytime, anywhere), mobile devices and access networks will have to work together in order to enable users to take full advantage of all available options.

The IEEE 802.21 working group (see www.ieee802.org/21) recently finalized the first standard for dealing with handovers in heterogeneous networks, also called *Media-Independent Handovers* (MIH)^[3]. The standard is expected to allow mobile users (and operators) to take full advantage of overlapping and diverse access networks. It provides a framework for efficiently discovering networks in range and executing intelligent heterogeneous handovers, based on their respective capabilities and current link conditions. This article aims to serve as a primer for those interested in the IEEE 802.21 standard. After introducing the IEEE 802.21 reference model, we present the MIH services and provide illustrative use cases that highlight the benefits of employing the Media-Independent Handover Services standard in heterogeneous networks.

Mobile and Wireless

The widespread success of 3G technologies^[4, 5] is evidenced by the rapid increase in the amount of data traffic over cellular networks in recent years. In Sweden, for example, the total amount of mobile data traffic leapt tenfold from just over 203 TB in 2006 to 2191 TB in 2007^[6]. This trend is expected to continue unabated with the deployment of *High-Speed Packet Access* (HSPA) and *Long-Term Evolution* (LTE) in the coming years. Of course, the amount of traffic over cellular networks is only a proportion of the traffic that originates from or terminates at WLANs worldwide. Campuswide deployments of WLANs are becoming the norm in developed countries, and we even find citywide WLANs, as in the case of the city of Oulu, Finland (see www.panoulu.net).

Finally, many anticipate that mobile WiMAX^[7] deployments will significantly affect telecommunications markets. In short, we are moving toward a far more heterogeneous network access environment than the one users and operators face today, with multiple overlapping mobile and wireless networks with diverse characteristics.

Multiaccess Devices in Heterogeneous Networks

As communication environments become more complex because of the diversity of network access technologies that support, for example, different access rates and *Quality of Service* (QoS) levels, users expect more from their wireless operator. Mobile devices, once featuring tiny screens, extremely limited processing and storage capacities, and narrowband connectivity^[8], now pack capabilities that just a few years ago were typical of high-end laptops. This scenario has allowed users to increasingly depend on mobile devices for e-mail and *Instant Messaging* (IM), but also for making *Voice over IP* (VoIP) calls, listening to streaming Internet radio, and watching online videos.

With respect to user mobility patterns, campuswide Wi-Fi users typically spend most of their connection time attached to a small set of access points located within a small radius^[9, 10]. This situation is not surprising, because Wi-Fi was originally designed and subsequently deployed mainly as an extension to wired infrastructures. In the future, however, we anticipate that multiaccess devices will employ different network interfaces to attach to different access networks, establishing multiple parallel connections over 3G/*Universal Mobile Telecommunications Service* (UMTS) and Wi-Fi, for example. With global reachability and ABC mechanisms in place, mobile devices will be able to selectively connect to different access networks depending on certain criteria. Keep in mind that from a conventional, IP-centered point of view, changing the *Point of Attachment* (PoA) calls for mobility management actions^[11, 12, 13], although in practice there may be no physical mobility whatsoever.

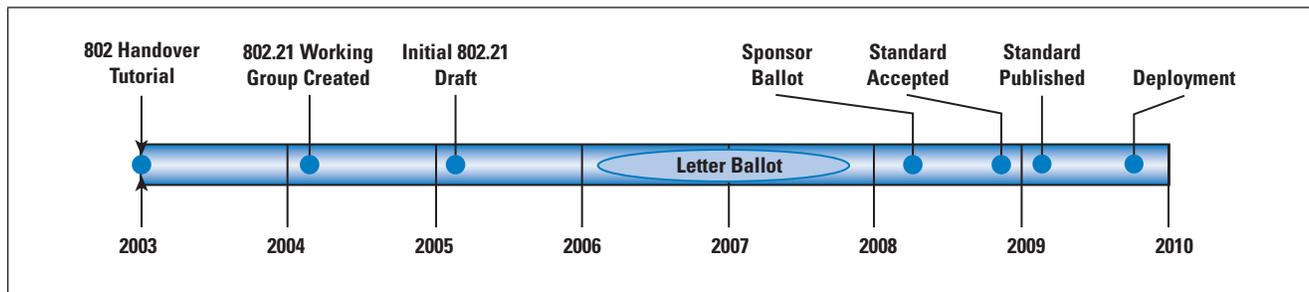
Given the diversity of networked applications running on mobile devices, knowledgeable network resource planning and operation is needed, in turn calling for a framework that allows users and their applications to state their network access preferences. This framework should also allow operators to steer terminal access patterns aiming at maximizing resource usage and increasing user satisfaction. For instance, podcasts can be downloaded only when connected to an uncongested WLAN, but web, map/navigation, and e-mail clients can use the cellular network or WLAN access on demand. Currently, this process can only be done manually: users need to be watchful for available access networks and choose which one to attach to based on very rudimentary information such as signal quality. If mobile nodes could collect timely and consistent information about the state of all available networks in range and were given the means to control their network connectivity, then a whole range of possibilities would become available.

In order to optimize the use of available network resources, mobile nodes need to be able to collect information about numerous heterogeneous networks in a generic and standardized way, irrespective of the underlying network access technology. The collected information, both dynamic and static, can then be used by handover decision-making processes, such as, say, mobility managers. Mobility managers can be enhanced versions of *Mobile IP* (MIP)^[11, 12, 13], proprietary solutions, or other proposals stemming from recent research, such as [14]. Researchers in the area have proposed several cross-layer frameworks for enhancing the efficiency of handover decision makers (see [14, 15] and the references therein), but none of them has been formally standardized or is widely accepted so far. What is needed is a standard framework that can attract ample support from major vendors and operators, and can be deployed incrementally.

Introducing IEEE 802.21-2008

Figure 1 illustrates the progress toward the IEEE 802.21-2008 standard. The working group was initiated in 2004, and the latest draft version of the standard was accepted as a new standard by the IEEE-SA Standards Board in November 2008^[3]. The standard was published in January 2009. It is anticipated that actual deployment of the standard will take place at the earliest in late 2009–2010.

Figure 1: Timeline of the IEEE 802.21-2008 Standardization Effort



IEEE 802.21-2008, also known as *Media-Independent Handover Services*, features a broad set of properties that meet the requirements of effective heterogeneous handovers. It allows for transparent service continuity during handovers by specifying mechanisms to gather and distribute information from various link types to a handover decision maker. The collected information comprises timely and consistent notifications about changes in link conditions and available access networks.

Note that the scope of IEEE 802.21-2008 is restricted to access technology-independent handovers. Intratechnology handovers, handover policies, security mechanisms, media-specific link layer enhancements to support IEEE 802.21-2008, and *Layer 3* (L3) and upper-layer enhancements are outside the scope of IEEE 802.21-2008. This article summarizes the salient points of [3], which henceforth is referred to as IEEE 802.21.

The IEEE 802.21 Reference Model

IEEE 802.21 facilitates a variety of handover methods, including both *hard handovers* and *soft handovers*. A hard handover, also known as “break-before-make” handover, typically implies an abrupt switch between two access points, base stations, or, generally speaking, PoAs. Soft handovers require the establishment of a connection with the target PoA while still routing traffic through the serving PoA. In soft (“make-before-break”) handovers, mobile nodes remain briefly connected with two PoAs. Note, however, that depending on service requirements and application traffic patterns, hard handovers may often go unnoticed. For example, web browsing and audio/video streaming with prebuffering can be accommodated when handing over between different PoAs in the range of one network by employing mechanisms that allow transferring the node connection context from one PoA to another quickly.

The main design elements of IEEE 802.21 can be classified into three categories: a framework for enabling transparent service continuity while handing over between heterogeneous access technologies; a set of handover-enabling functions; and a set of *Service Access Points* (SAPs).

Transparent Service Continuity

IEEE 802.21 specifies a framework that enables transparent service continuity while a mobile node switches between heterogeneous access technologies. The consequences of a particular handover need to be communicated and considered early in the process and, clearly, before the handover execution. In soft handovers, it is crucial that service continuity, during and after the handover, is ensured without any user intervention. To this end, IEEE 802.21 specifies essential mechanisms to gather all necessary information required for an affiliation with a new access point before breaking up the currently used connection. Interactive applications, such as VoIP, are typically the most demanding in terms of handover delays, and high-quality VoIP calls can be served only by soft handovers. On the other hand, video streaming can accommodate hard handovers, as long as the vertical break-before-make handover delay does not exceed the application buffer interval delay. In the case of hard handovers, handover preparation signaling can initiate the connection context transfer from the serving PoA to the target PoA beforehand.

For instance, lack of the required level of QoS support or low available capacity in a candidate access network may lead the network selecting entity to prevent a planned handover. On the other hand, for example, increasing delay, jitter, or packet-loss rates in the currently serving network may degrade the perceived QoS throughout the network, or only for a particular application, triggering the mobility manager to start assessing the potential of candidate target access networks and subsequently initiate an IEEE 802.21-assisted handover.

IEEE 802.21 also allows the reception of dynamic information about the performance of the serving network and other networks in range. In other words, IEEE 802.21 provides methods for continuous monitoring of available access conditions. However, IEEE 802.21 does not specify any methods for collecting this dynamic information at the link layer.

Handover-Enabling Functions

IEEE 802.21 defines a set of handover-enabling functions, which are specified with respect to existing network elements in the protocol stack, and introduces a new logical entity called *Media-Independent Handover Function* (MIHF). The MIHF logically resides between the link layer and the network layer. It provides, among others, abstracted services to entities residing at the network layer and above, called *MIH Users* (MIHUs). MIHUs are anticipated to make handover and link-selection decisions based on their internal policies, context, and the information received from the MIHF. To this end, the primary role of the MIHF is to assist in handovers and handover decision making by providing all necessary information to the network selector or mobility management entities. The latter are responsible for handover decisions regardless of the entity position in the network. The MIHF is not meant to make any decisions with respect to network selection.

Service Access Points

SAPs with associated primitives between the MIHF and MIHUs (MIH_SAP) give MIHUs access to the following services that the MIHF provides:

- The *Media-Independent Event Service* (MIES) provides event reporting about, for example, dynamic changes in link conditions, link status, and link quality. Events can be both local and remote. Remote events are obtained from a peer MIHF entity.
- The *Media-Independent Command Service* (MICS) enables MIHUs to manage and control the parameters related to link behavior and handovers. MICS provides a set of commands for accomplishing that, as we will see later in this article. Commands can be both local and remote. The information obtained with MICS is dynamic.
- The *Media-Independent Information Service* (MIIS) allows MIHUs to receive static information about the characteristics and services of the serving network and other available networks in range. This information can be used to assist in making a decision about which handover target to choose and to make preliminary preparations for a handover.

Figure 2 illustrates the general reference model of IEEE 802.21. The scope of IEEE 802.21 includes only the operation of MIHF and the primitives associated with the interfaces between MIHF and other entities. A single media-independent interface between MIHF and MIHU (MIH_SAP) is sufficient.

On the other hand, there is a need for defining a separate technology-dependent interface, which is specific to the corresponding media type supported, between the MIHF and the lower layers (MIH_LINK_SAP).

The primitives associated with the MIH_LINK_SAP enable MIHF to receive timely and consistent link information and control link operation during handovers. For example, the currently supported link layers include wired and wireless media types from the IEEE family of standards (for example, 802.3, 802.11, 802.15, and 802.16), as well as those defined by the *Third-Generation Partnership Project* (3GPP) and *Third-Generation Partnership Project 2* (3GPP2). Besides these, IEEE 802.21 specifies a media-independent SAP (MIH_NET_SAP), which provides transport services for Layer 2 (L2) and Layer 3 (L3) MIH message exchange with remote MIHFs. Functions over the LLC_SAP are not specified in IEEE 802.21.

Figure 2: The IEEE 802.21-2008 Reference Model

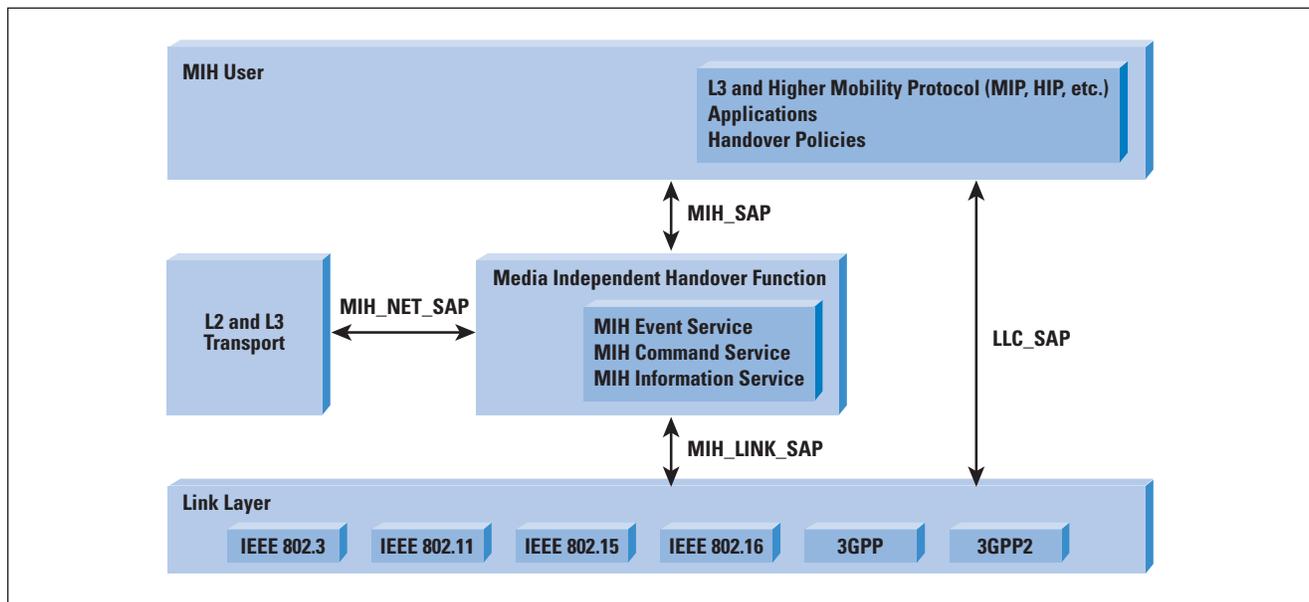
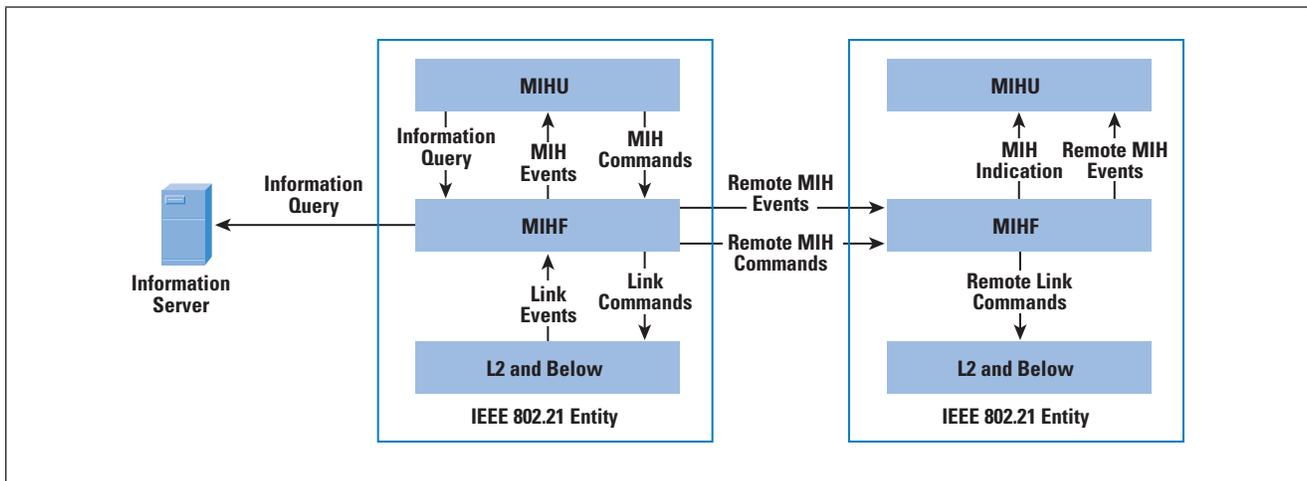


Figure 3 presents the messages directions of each MIHF service class, including both local and remote events and commands. The MIHF can subscribe to particular sets of events from a peer MIHF. Remote commands are initiated by local MIHUs and are conveyed to the peer MIHF through the local MIHF. Finally, MIIS information can be obtained through queries to the local database and to remote Information Servers.

Figure 3: MIHF Services



IEEE 802.21 Illustrated

Figure 4 illustrates an example topology where different wireless networks overlap. Imagine that the multiaccess mobile device user watches a high-bitrate IPTV channel as she moves in this area. Three wireless access technologies are considered in this example: Wi-Fi (IEEE 802.11), WiMAX (IEEE 802.16), and 3G/UMTS (3GPP). In this example, we assume that all networks and the mobile device are IEEE 802.21-compatible and that the Wi-Fi area is covered by several 802.11 PoAs, as would be the case in a campus- or citywide deployment.

Figure 4: Example Topology with Heterogeneous Overlapping Wireless Access Networks

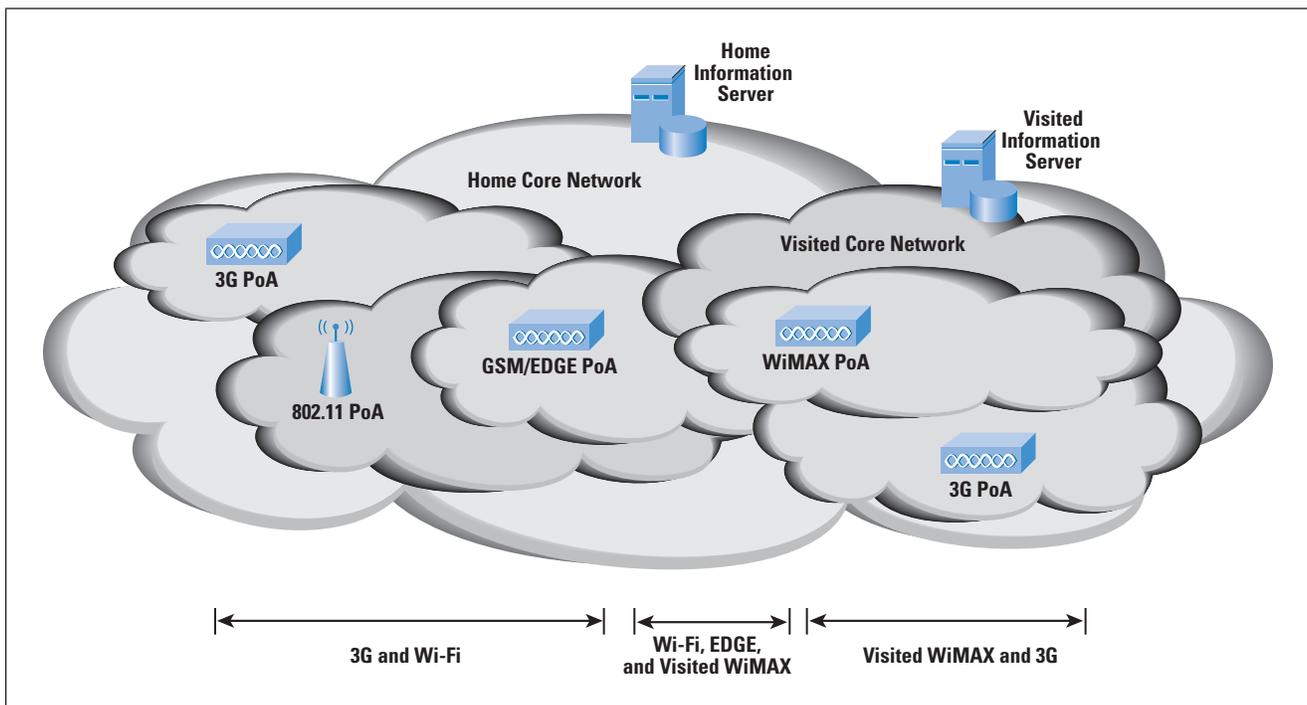
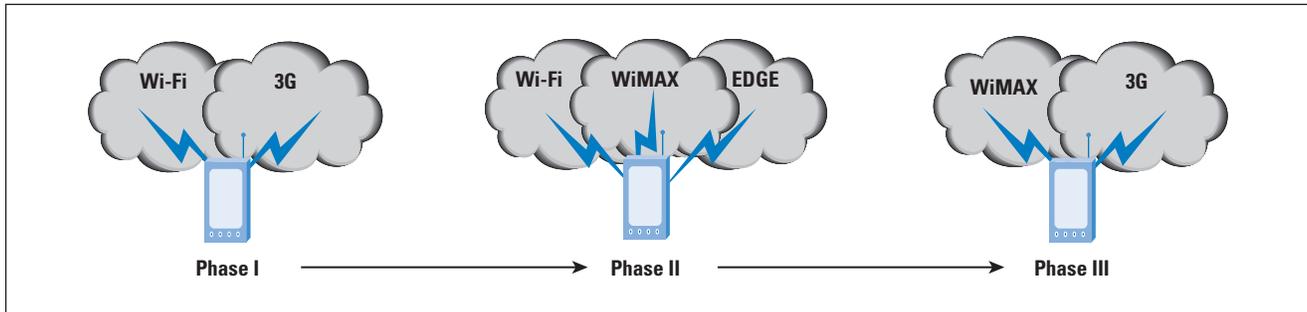


Figure 5 illustrates the network access environment as perceived by a mobile device in the area. The figure depicts three snapshots, indicating the overlapping networks in range at different locations. In order to deliver the IPTV stream transparently, for each of the available access networks we need to consider their effective available bandwidth, the associated cost per traffic unit, the terminal speed, the cell coverage area, the level of QoS support it can provide, and so on. Using information made available through the MIHF, we can determine which should be the next target access network.

Figure 5: Example Network Environment in Different Locations



In Phase I, the mobile node has two network access options. It can use a free and open Wi-Fi network or connect to the cellular operator's 3G/UMTS network. Note that opting to use the latter may, for instance, depend on the charging scheme of the operator. If subscribers pay based on traffic volume, one would assume that the free Wi-Fi network is a better option. On the other hand, as flat-rate plans become more popular, 3G may be a better option with its extended coverage and QoS guarantees. The IEEE 802.21 MIIS can provide this type of information, allowing for automation in dynamic access selection.

In Phase II, as the user moves, the device goes through a cellular technology handover from 3G/UMTS to *Enhanced Data rates for GSM Evolution* (EDGE)^[8]. At the same place, the public Wi-Fi network is still available and a new WiMAX network has just been detected. Assume that EDGE is not sufficient for delivering the IPTV stream. If in Phase I the network selection process opted for using the cellular network, then in Phase II the client application will experience significant degradation in service if it continues to use the EDGE access network. A vertical handover to the Wi-Fi or the WiMAX network should be considered. In contrast, if the mobile node first chose to stream the IPTV channel over the Wi-Fi access network, then it may need to reassess the situation based on events and link parameter reports using MIES and MICS, as we explain in the following sections. For example, an information query can reveal whether the WiMAX network is operated by a partner *Internet Service Provider* (ISP), and what the roaming cost would be.

Finally, in Phase III, the coverage area of the public Wi-Fi network ends. Through IEEE 802.21 services we find out that the only available networks are the roaming partner WiMAX and the home cellular network that is now offering 3G service.

The environment with several overlapping networks described previously and illustrated in Figures 4 and 5 is already a reality today in many places, and it is widely anticipated to be prevalent in the future. Next, we examine the three services defined by IEEE 802.21, namely MIES, MICS, and MIIS.

Media-Independent Event Service

Events indicate or predict changes in the state and transmission behavior of physical, data link, and logical link layers. In general, events are triggers for initiating candidate network discovery and handover procedures. The events defined in IEEE 802.21 are categorized as either *Link Events* or *MIH Events*, depending on their origin. Link events emanate from the link layers, whereas MIH events emanate from the MIHF and can be both remote and local. Local events propagate from lower layers to upper layers through the MIHF. Remote events occur at the protocol stack of another network entity and are transmitted from a peer MIHF to the local MIHF, as illustrated in Figure 3.

The *Media-Independent Event Service* (MIES) currently supports five types of events: MAC and PHY State Change events, Link Parameter events, Predictive events, Link Handover events, and Link Transmission events. A short introduction to the event types and corresponding events follows.

MAC and PHY State Change events correspond to state changes in MAC and *physical* (PHY) layers. The most characteristic events in this category are *Link_Up* and *Link_Down* events, which are generated when a Layer 2 connection with an access point is established or is torn down, respectively. Another event, called *Link_Detected*, indicates that a PoA has been detected but no affiliation is established yet.

Link Parameter events relate to changes in Layer 2 parameters. A *Link_Parameters_Report* can be sent when a MIHU has set thresholds for certain parameters. For example, a MIHU can set thresholds for the *Received Signal Strength Indicator* (RSSI) on IEEE 802.11 links, so that when a threshold is crossed proper action can be taken. A *Link_Parameters_Report* is also used for issuing periodical notifications about link conditions. Based on Link Parameter events, a MIHU can initiate the handover candidate discovery process, or trigger applications to adapt to changing link conditions.

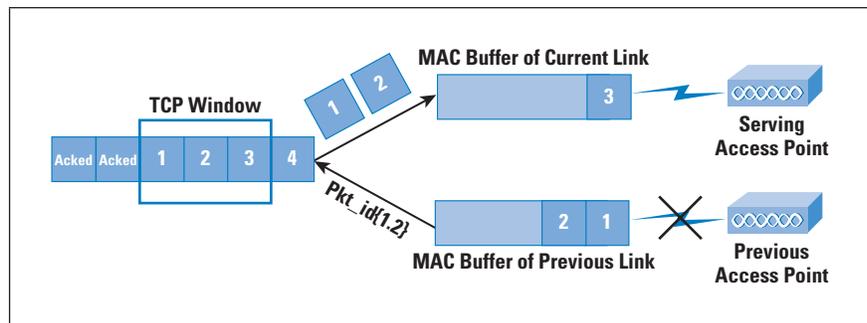
Predictive events inform about the probability of dramatic (negative) changes in link characteristics in the near future. For example, if strong decay in signal strength is observed, this decay may indicate imminent loss of link connectivity. Predictive events may include temporal information about when the actual event is expected to occur and what its presumed likelihood is. A *Link_Going_Down* event, for instance, may trigger a MIHU to consider possibilities for handing over to other available networks in range.

Link Handover events indicate the occurrence of Layer 2 handovers. The *Link_Handover_Imminent* event serves as a notification for an imminent handover, whereas a *Link_Handover_Complete* event reports the successful change of PoA. These events emanate from the link layer and are based solely on local Layer 2 information.

Link Transmission events show the transmission status of individual higher-layer *Protocol Data Units* (PDUs) at the link layer. Upper layers can, for example, adapt to data loss during a handover by improving buffer management based on Link Transmission events. These events may allow future upper-layer implementations to identify lost packets and recover without waiting for the expiration of retransmission timers.

Currently, for example, in the case of an ongoing session over TCP, the occurrence of a handover may have dramatic effects in performance. With IEEE 802.21, MIHUs can be informed about individual packets that have already been delivered to the sending buffer of the MAC layer but were not successfully transmitted before the handover occurred. In other words, the MAC layer outgoing buffer may contain TCP segments that cannot be delivered through the wireless network to the peer at the other end of the TCP connection. These segments were not successfully delivered from the local *Automatic Repeat-reQuest* (ARQ) module over the first hop, but are still buffered and cannot be transmitted because there is no link connectivity. In this case, TCP could use the information from Link Transmission events that identifies which packets need to be resent through the new access network, as illustrated in Figure 6 for packet numbers 1 and 2. Note, however, that IEEE 802.21 does not define any identifier for reliable packet identification, only the size of the packet ID (2 bytes), and it is up to the implementer to determine how different messages will be locally identified.

Figure 6: Link Transmission Event Indicating Undelivered Packets



Media-Independent Command Service

The *Media-Independent Command Service* (MICS) enables higher layers to control the stream of events originating from lower layers. Commands can originate from MIHUs (MIH commands) or from the MIHF (Link commands) and the destination can be the MIHF or any lower layer, respectively, as shown in Figure 3. The responses to Link commands are sent to MIHUs as indications. MIHUs can use command services to determine the status of different links in a uniform way, and control each interface accordingly, aiming for optimal connectivity. MICS defines the following set of commands that enable MIHUs to configure, control, and get information from the lower layers:

- *MIH commands* can be directed to lower layers residing at both local and remote MIHF entities. They originate from the upper layers and are directed to the MIHF. Similarly with MIH events, MIH commands can be both remote and local. MIH commands are typically used for network selection and handover management because they allow upper layers to initialize, prepare for, and execute handovers. MIH commands are also used to configure custom thresholds for link parameters. As mentioned previously, when set thresholds are crossed, MIHUs get the corresponding notifications through Link Parameter events.
- *Link commands* originate from the MIHF and are sent to lower layers in order to control their operation. Link commands can be issued only locally. Nevertheless, Link commands can be executed on behalf of local MIHUs, which could act on information received from a remote peer. Link commands are often initiated by MIHUs. For example, an MIHU can issue the *MIH_Get_Link_Parameters* MIH command, which when received by the local MIHF will lead to the generation of a remote *Link_Get_Parameters* Link command, as shown in Figure 3. This way, the MIHF can acquire the current parameter values of active link(s) for MIHU, and then deliver this information to the requesting MIHU. Note that MICS provides dynamic information about different link parameters, in contrast with MIIS, described next, which can report only static information.

Media-Independent Information Service

The *Media-Independent Information Service* (MIIS) facilitates handovers through a unified set of mechanisms that the MIHF can use to discover and obtain static (or rarely changing) information about networks in the vicinity of a multiaccess node. In other words, MIIS allows mobile nodes to check for available networks in range while using their currently active access network. MIIS information exchange occurs at the link layer (Layer 2) or network layer (Layer 3), so that all necessary information related to link layer or higher-layer services is collected before a mobile node authenticates with a new PoA.

MIIS defines a set of *Information Elements* (IEs) that are indispensable for network selection, classified into three groups: General Information and Access Network-Specific Information; PoA-Specific Information; and Other Information, which includes vendor- and network-specific details. The types of information handled by MIIS are solely related to handover decisions and conformance to the affiliation with the new PoA. Information relevant for assessing candidate networks by the handover machinery includes connection establishment details, such as PoA address and location; which security mechanisms are supported in a given access network; and what QoS guarantees can be provided.

General Information Elements and *Access Network-Specific Information Elements* give a general overview of neighboring networks. Information Elements may include, for instance, a list of available networks and their associated operators, roaming agreements and costs, and security and QoS support. For instance, user policies, defined at higher layers, may dictate that if a given access network operator charges users based on their traffic volume, then the network selector entity should not consider the corresponding access when a high-bitrate service, such as IPTV, is active.

PoA-Specific Information Elements refer to each PoA available in the access network and report PoA location and addressing information, supported data rates, PHY and MAC layer types, and channel parameters that can optimize link layer connectivity. Some additional information related to higher-layer services and individual capabilities of particular PoAs may be included as well. For instance, an advanced mobility manager on the mobile node can use the information about the geographical position of a PoA and compare it with the current or expected node location based on its mobility patterns. With careful planning and by taking advantage of this information, mobile nodes may be able to reduce the number of handovers and optimize the use of network resources.

MIIS provides mechanisms for issuing and responding to queries for Information Elements. Such information may reside in a separate server or in a local information database at the mobile node (see Figure 3). An MIHF could have access to an information server in its IEEE 802.21-enabled *Point-of-Service* (PoS) range from which it can obtain information regarding the home PoS and possibly other PoSs, such as those of roaming partners. If the home information server is not able to provide any information regarding the visited network, an MIIS query can be directed to the peer MIHF, residing in the visited PoS, which can access the visited PoS information server. Information queries can often be answered locally, based on information gathered from previous queries and by preprovisioning, for example, from the information server.

Information Elements and their relationships are captured in an Information Service schema which, in turn, defines the information structure. IEEE 802.21 specifies that information that is to be presented across different technologies should be in a standardized, common, and open format, such as XML or *Type Length Value* (TLV).

Service Management

In order to use and provide MIHF services, MIHF entities need to be configured appropriately. IEEE 802.21 defines three service management functions: MIH capability discovery, MIH registration, and MIH event subscription.

MIHF may discover other MIHF entities and their capabilities using the MIH capability discovery procedure. Depending on the information obtained from this procedure, the local MIHF can determine which peer MIHFs it should register with. The MIH capability discovery function uses the MIH protocol (introduced in the following section) at Layer 2 or Layer 3, and media-specific Layer 2 broadcast messages are allowed. For example, an MIHF can listen to media-specific broadcast messages, such as IEEE 802.11 beacons, or media-independent Layer 2 *MIH_Capability_Discover* broadcast messages, because an MIHF entity residing in the network may announce its existence and capabilities periodically. MIHF can also send *MIH_Capability_Discover* request messages using multicast or unicast to detect peer MIHFs in a solicited way. For instance, MIHF can send a request by unicast for obtaining the capabilities of a specific IEEE 802.21 network entity. In this case, only the IEEE 802.21 network entity addressed should respond to these request messages.

MIH registration is a symmetric procedure by which two peer MIHFs authenticate and can then communicate with each other in a more trusted manner. After MIH registration is completed, the two peer MIHF entities can symmetrically request services from their registered peer. Note that MIH registration is not necessary for obtaining some level of support from a peer MIHF. However, by registering and authenticating, peer MIHFs typically will get access to much more extensive information. That is, although the MIHF residing on the mobile node may be able to access information services from the network-side MIHFs without registration and authentication, the available information may be only a subset of that provided after authenticating.

Finally, MIH event subscription enables MIHUs to subscribe to a particular set of events provided by MIES from the local or peer MIHF. Event subscription from a peer MIHF requires registration and knowledge about its capabilities. The subscription contains only the list of events the MIHU is interested in. Note that event sources may not be necessarily capable of providing all events that the subscriber is interested in subscribing to. Each subscription request is matched by a confirmation message from the event source indicating the events approved for subscription.

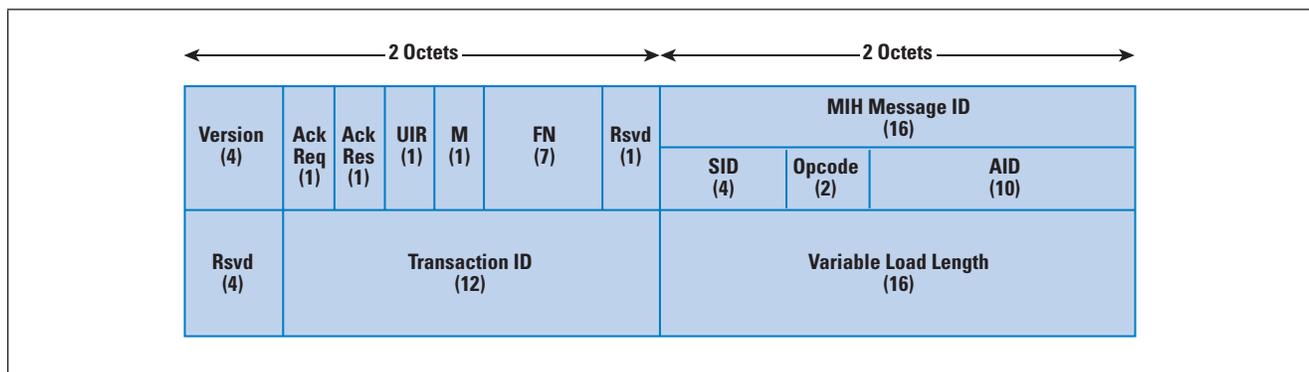
Media-Independent Handover Protocol

The *Media-Independent Handover Protocol* (MIHP) specifies the rules and services for unified communication between peer MIHFs. The protocol defines the message format, header, and encoding format and is meant to be used solely for communicating with peer MIHF entities. For internal communication no particular encoding is dictated.

MIH protocol messages can be carried over Layer 2 management frames, Layer 2 data frames, or over Layer 3/IP transport. Note that cellular technologies do not provide Layer 2 transport without changes in their protocol stack.

The MIH protocol messages, or frames, comprise a header part and a TLV-encoded payload part. The MIHF frame header consists of eight octets. Figure 7 illustrates the MIH protocol header indicating the corresponding bit length for each field in parentheses.

Figure 7: MIH Protocol Header



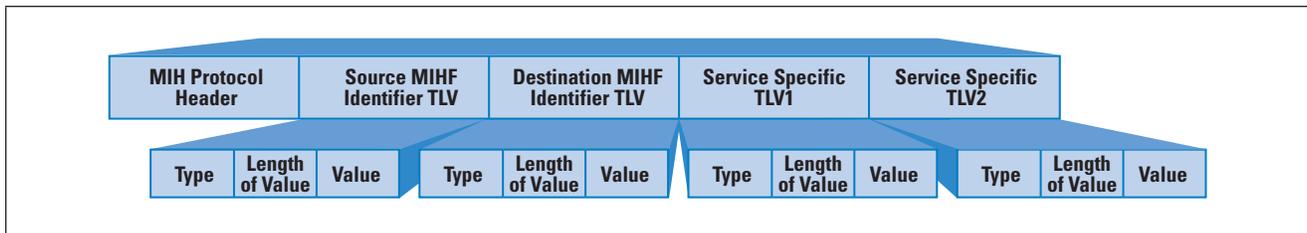
The *Version* field in the MIH frame header specifies the version of the MIH protocol used. The two *Ack* fields are for acknowledgement purposes and are discussed later in the article. The *Unauthenticated Information Request* (UIR) flag indicates that the response message may be sent with a limited length because of the nature of unauthenticated message exchange. Recall that when an MIHF issues requests without registering first with its peer, it may receive less information than if it had registered earlier. If this flag is set, then the information included in the response message may not reflect the complete information available to registered MIHFs. The *More Fragments* (M) and *Fragment Number* (FN) fields are used in message fragmentation.

The *MIH Message ID* field comprises three subfields. The *Service Identifier* (SID) field indicates the MIHF service class (MIES, MICS, MIIS, or Service Management) that this message belongs to. The *Operation code* (Opcode) specifies whether the message is a request, response, or indication. The *Action Identifier* (AID) is related with and scoped by the SID. For instance, if the SID indicates MIES, AID points to the actual event type. The *Variable Load Length* field contains the total length of the variable, TLV-encoded payload carried by this message frame.

The MIH protocol messages use the *Transaction ID* and *MIHF ID* fields as identifiers, but only the former is included in the header. The Transaction ID field is an identifier that helps to match each request, response, or indication message with its acknowledgement.

The payload part contains service-specific messages encoded in TLV format. The first two TLVs in the payload part (not shown in Figure 7) should be the *Source Identifier* and *Destination Identifier*, which are both the same data type as the MIHF ID. Every MIHF must have a unique MIHF ID, which may be assigned to it at configuration time. The MIHF ID shall be invariant and could be, for example, a *Fully Qualified Domain Name* (FQDN) or *Network Access Identifier* (NAI). The MIHF ID is used during the MIH registration phase and is appended to the payload part of every message requiring endpoint identification. In broadcast messages, the Destination Identifier TLV is defined as zero length. Figure 8 shows the message structure consisting of the MIH Protocol header, source and destination identifiers, and service-specific TLVs. In TLV encoding, the Type field (1 octet) denotes the parameter type, the Length field (variable octets) indicates the length of the Value field, and the Value field (variable octets) carries the actual value of the parameter.

Figure 8: MIH Protocol Frame Structure



Acknowledging MIH messages is not mandatory. Still, the MIH protocol does support the use of acknowledgements to ensure reliable message exchange. The sender MIHF can set the *ACK-Req* field to instruct the receiver to return an acknowledgement with *ACK-Rsp* bit set. The *MIH Message ID* and *Transaction ID* must be the same in the request message and its acknowledgement. An acknowledgement message may carry no payload. Note, however, that despite employing these two ID fields, the MIH protocol does not specify any further mechanisms for reliable authentication or shielding message exchanges from third parties.

MIH Communication Model

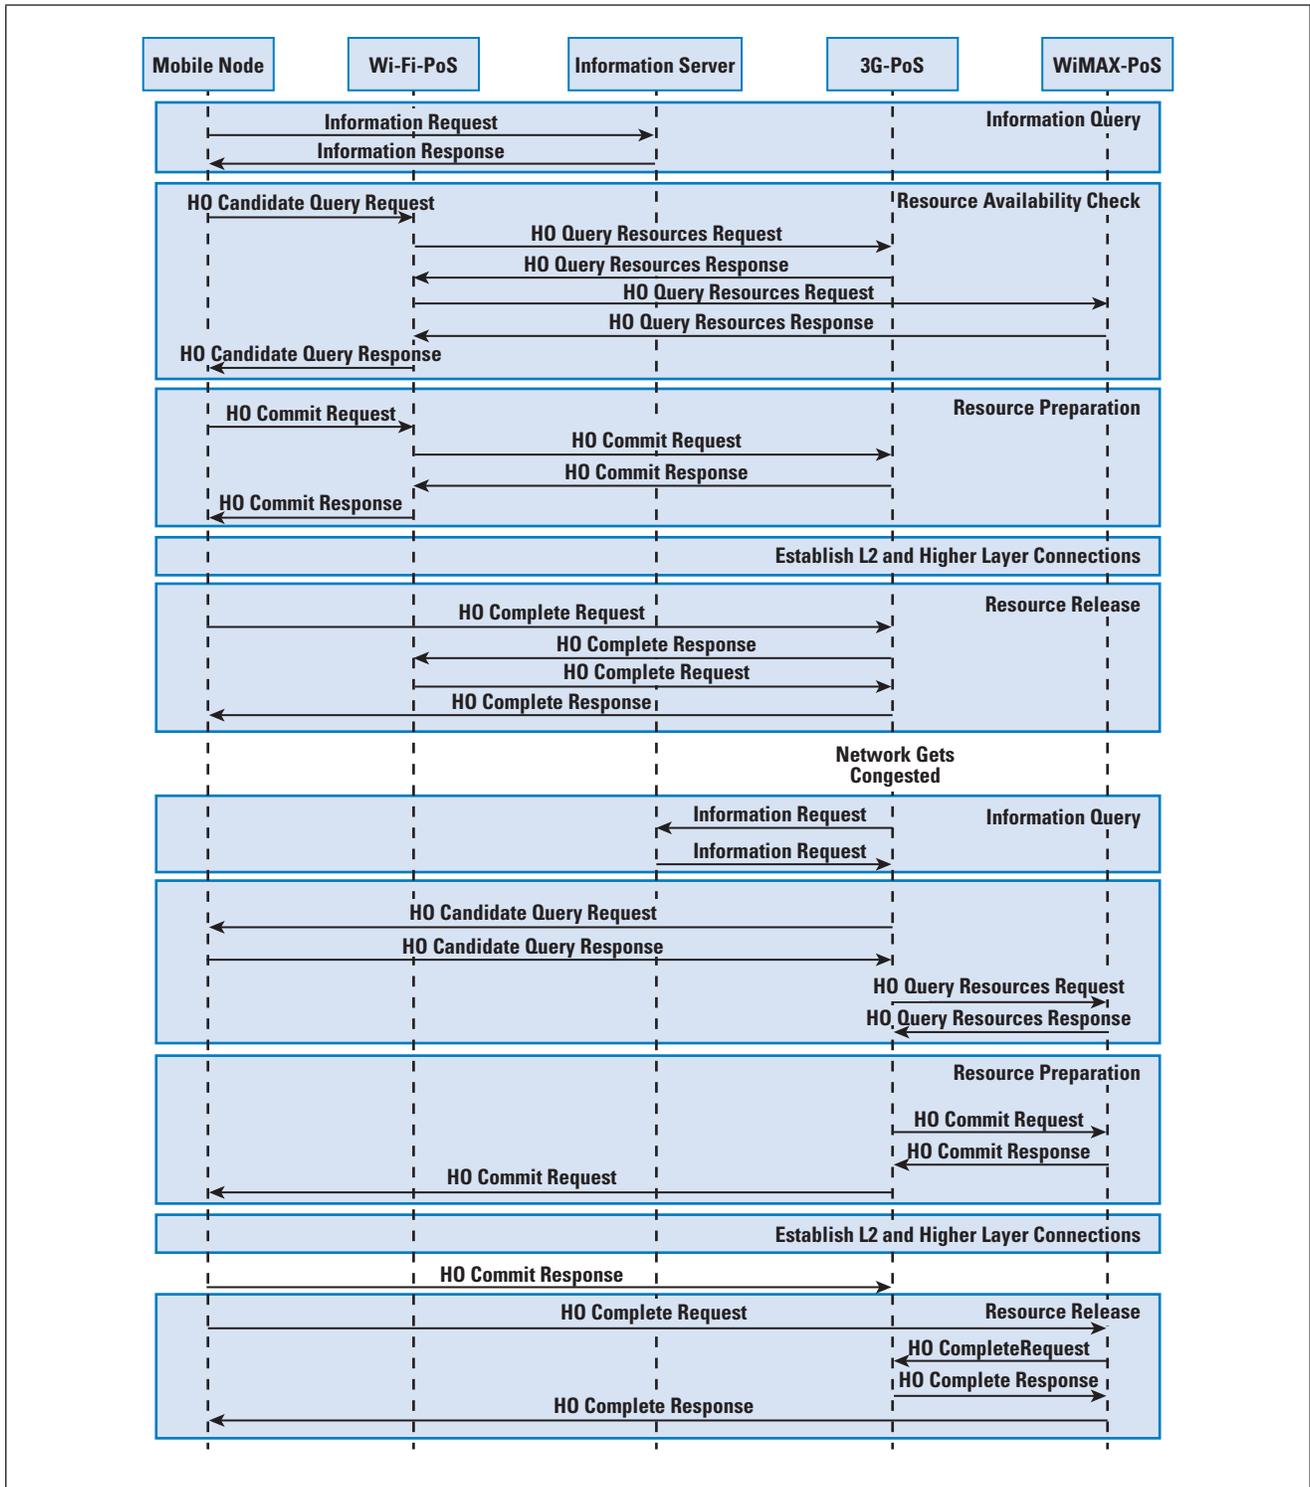
The MIHF communication model specifies different MIHF roles and their communication relationships, such as supported transport mechanisms and service classes. The assigned MIHF roles depend on their location in the network. For example, an MIHF on a mobile node can communicate directly with network-side entities called *MIH PoSs* using Layer 2 or Layer 3 communication. MIH PoSs may include the serving PoA or candidate PoAs. Network-side MIHFs can communicate with each other at Layer 3 or above using the MIH protocol, introduced in the previous section.

Let us revisit the example use case of IEEE 802.21 illustrated in Figures 4 and 5. Figure 9 presents the IEEE 802.21 message exchanges in mobile- and network-initiated handover procedures in the case where the mobile node hands over from a Wi-Fi to the 3G cellular network (between Phase II and Phase III in Figure 5) and then hands over to a WiMAX network (Phase III in Figure 5). First, during the discovery of handover candidate PoAs, the mobile node MIHF employs MIIS to gather static information about the surrounding networks. The request is issued over the currently used Wi-Fi access. This information is obtained from the information server that may reside in a different network than the one currently in use.

After receiving the response to its Information Request, the mobile node initiates the handover process by querying about the availability of resources in the networks it is interested in. These requests are sent through the serving PoS (*Wi-Fi-PoS* in Figure 9), which disseminates the requests to the MIH PoSs of the candidate networks (*3G-PoS* and *WiMAX-PoS* in Figure 9). The response indicating the capabilities of the two candidate networks is returned to the mobile node MIHF from the serving PoS. After receiving this information, an MIHU on the mobile node decides which network to hand over to, based on policies and the output of its network selection algorithms. Then a *Handover Commit Request* message is sent, and after the candidate network has made its final commitment for the handover (and the appropriate resources are reserved successfully), the mobile node establishes a Layer 2 connection with the PoA in the area of the candidate PoS, that is, the *3G-PoS* in our example case. Following this successful intertechnology handover, the resources used in the previous link can optionally be released. In the case where no resources are explicitly reserved, this step is skipped.

As we progress in the timeline of our example case, the network-side MIHU initiates a handover to the WiMAX network. This handover could be, for example, the result of observing congestion in the cellular network that indicates that a new PoS should be found for the mobile node. The serving PoS (*3G-PoS*) collects information about networks in the range of the mobile node from the Information Server. Upon determining that a suitable WiMAX candidate network that can serve the mobile node exists, the *3G-PoS* triggers a network-initiated handover. First, the serving PoS requests permission from the mobile node to proceed with the handover. If the mobile node does not object, the serving PoS proceeds with the rest of the handover procedure, which is similar to the mobile-initiated handover described previously except that it is handled by a network entity.

Figure 9: IEEE 802.21-Assisted Handover Message Sequence Diagram



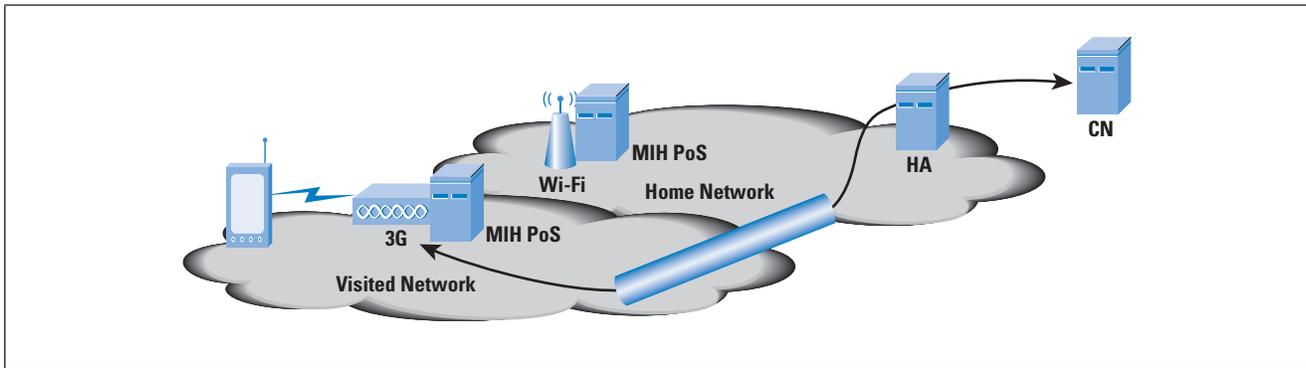
Handover Execution

As illustrated in the example, the handover decision and target assessment constitute a multiphase process where the assistance of IEEE 802.21 is essential. However, the actual handover execution is outside the scope of the standard. This section briefly describes how handovers can be carried out by MIP with the cooperation of IEEE 802.21. After choosing the target network by capitalizing on the IEEE 802.21 services, the mobile node establishes a new connection with the handover target network while still routing traffic through the currently serving network. The mobile node obtains a *Care-of Address (CoA)* for this new link from the IP address space of the target network. The CoA is an IP address assigned to the new link of the mobile node and is used while connected to the visiting network^[11]. With MIPv4, the CoA is provided by a *Foreign Agent (FA)* in the visited network, which also acts as a router for the mobile node^[12]. With MIPv6, the Foreign Agent is not needed^[13] and the CoA is obtained directly, say, for example, from a *Dynamic Host Configuration Protocol (DHCP)* sever. The mobile node can obtain the IP address of the DHCP server in the target network through the IEEE 802.21MIIS.

In MIP, each mobile node has a *Home Agent (HA)*, which routes the traffic of the mobile node. After successfully affiliating with a PoA in the target network, the mobile node notifies the Home Agent of the CoA by performing a binding update. In a bidirectional tunnel mode, the Home Agent establishes an IP-IP tunnel between the Home Agent and the Foreign Agent (MIPv4) or the Home Agent and the mobile node CoA (MIPv6). This mode does not require any binding updates on the *Correspondent Node (CN)*. In other modes, either the uplink traffic of the mobile node is sent directly to the Correspondent Node using the CoA as source address, or all bidirectional communication between the Correspondent Node and the mobile node uses the CoA only. In the first case, traffic from the Correspondent Node to the mobile node travels through the Home Agent, but in the latter case there is no need for the Home Agent detour. However, these modes need address binding at the Correspondent Node and are in practice less frequently used than the bidirectional tunnel mode.

Figure 10 illustrates a situation where a link with the Wi-Fi PoA is broken down by the mobile node and the IPv6 traffic between the Correspondent Node and the mobile node, now employing IEEE 802.21-enabled 3G network, travels through the tunnel between Home Agent and the mobile node.

Figure 10: Mobile IPv6 Tunnel



Layer 3 handover executions based on RFC 3344^[12] and RFC 3775^[13] may often exceed the typical handover delay budgets, thus introducing gaps in connectivity that are perceptible at the application layer. Recent standardization efforts have focused on decreasing handover delays by enhancing MIP so that it can provide for transparent mobility management for both IPv4^[16] and IPv6^[17, 18]. The proposed enhancements either reduce the amount of signaling or allow the mobile node to configure the new Layer 3 connection before re-associating with the new network. In this context, IEEE 802.21 can provide the essential information for preestablishing the connection based on media-independent Layer 2 link detection events as well as static address information from the target network.

Summary and Outlook

We presented an overview of the IEEE 802.21 Media-Independent Handover Services standard. We anticipate that its adoption in the near future will allow for better network resource usage and permit multiaccess devices to select the network access best suited for their communication needs. After motivating the needs for a standard to cope with heterogeneous network handovers, we introduced the IEEE 802.21 Reference Model and the MIH Services. We briefly presented the MIH Protocol, although a more thorough description calls for a separate overview article. Finally, we illustrated network operation when IEEE 802.21 is adopted using example use cases featuring both network- and terminal-initiated intertechnology (or vertical) handovers.

We expect that in the future, when IEEE 802.21-2008 is widely deployed, there will be significant efforts to further amend and extend it in order to provide for even better services. In fact, because security mechanisms are outside the scope of the base IEEE 802.21 standard, the work on defining a security-related extension to IEEE 802.21 (IEEE P802.21a) has already begun. Moreover, another amendment (IEEE P802.21b) that deals with handovers with downlink-only technologies, such as *Digital Video Broadcasting* (DVB), has also been introduced (see www.ieee802.org/21 for more information about the amendments). Nevertheless, it remains uncertain whether vendors will stand by this promising standard and incorporate it in future products and solutions.

References

- [1] T. Sridhar, “Wi-Fi, Bluetooth and WiMAX,” *The Internet Protocol Journal*, Volume 11, No. 4, December 2008.
- [2] E. Gustafsson and A. Jonsson, “Always Best Connected,” *IEEE Wireless Communications*, Volume 10, No. 1, February 2003.
- [3] IEEE Std 802.21-2008, *IEEE Standard for Local and Metropolitan Area Networks—Part 21: Media Independent Handover Services*, IEEE, January 2009.
- [4] H. Kaaranen, S. Naghian, L. Laitinen, A. Ahtiainen, and V. Niemi, *UMTS Networks: Architecture, Mobility and Services*, 2nd Edition, John Wiley & Sons, 2005.
- [5] V. Vanghi, A. Damnjanovic, and B. Vojcic, *The cdma2000 System for Mobile Communications: 3G Wireless Evolution*, Prentice Hall, 2004.
- [6] Y. Mälärstig, O. Holmström, and P. Davidsson, *Svensk telemarknad 2007*, PTS-ER-2008:15, ISSN 1650-9862, June 2008.
- [7] J. Pinola and K. Pentikousis, “Mobile WiMAX,” *The Internet Protocol Journal*, Volume 11, No. 2, June 2008.
- [8] K. Pentikousis, “Wireless Data Networks,” *The Internet Protocol Journal*, Volume 8, No. 1, March 2005.
- [9] M. Balazinska and P. Castro, “Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network,” *Proc. First International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Francisco, California, USA, May 2003, pp. 303–316.
- [10] T. Henderson, D. Kotz, and I. Abyzov, “The Changing Usage of a Mature Campus-wide Wireless Network,” *Computer Networks*, Volume 52, No. 14, October 2008, pp. 2690–2712.
- [11] W. Stallings, “Mobile IP,” *The Internet Protocol Journal*, Volume 4, No. 2, June 2001.
- [12] C. Perkins (Ed.), “IP Mobility Support for IPv4,” RFC 3344, August 2002.
- [13] D. Johnson, C. Perkins, and J. Arkko, “Mobility Support in IPv6,” RFC 3775, June 2004.

- [14] K. Pentikousis, R. Agüero, J. Gebert, J. A. Galache, O. Blume, and P. Pääkkönen, “The Ambient Networks Heterogeneous Access Selection Architecture,” *Proc. First Ambient Networks Workshop on Mobility, Multiaccess, and Network Management (M2NM)*, Sydney, Australia, October 2007, pp. 49–54.
- [15] J. Mäkelä and K. Pentikousis, “Trigger Management Mechanisms,” *Proc. Second International Symposium on Wireless Pervasive Computing (ISWPC)*, San Juan, Puerto Rico, February 2007, pp. 378–383.
- [16] K. El Malki (Ed.), “Low Latency Handoffs in Mobile IPv4,” RFC 4881, June 2007.
- [17] H. Soliman, C. Castelluccia, K. El Malki, and L. Bellier, “Hierarchical Mobile IPv6 Mobility Management,” RFC 4140, August 2005.
- [18] R. Koodli (Ed.), “Mobile IPv6 Fast Handovers,” RFC 4068, July 2005.

ESA PIRI received his M.Sc. from the University of Oulu, Oulu, Finland, in 2008. During his studies, he specialized in information networks systems and wrote his Master’s thesis on mobility management issues in heterogeneous networks. Currently he is working as a Research Scientist at VTT Technical Research Centre of Finland in Oulu, Finland. He can be contacted by e-mail at: esa.piri@vtt.fi

KOSTAS PENTIKOUSIS studied computer science at Aristotle University of Thessaloniki, Greece (B.Sc. 1996, summa cum laude) and State University of New York at Stony Brook, USA (M.Sc. 2000, Ph.D. 2004). He has published internationally in several areas, including mobile computing; transport protocols; applications; network traffic measurements and analysis; and simulation and modeling. Dr. Pentikousis is a Senior Research Scientist at VTT Technical Research Centre of Finland. Visit <http://ipv6.willab.fi/kostas> for more information and contact details.

Book Review

Geeks Bearing Gifts

Geeks Bearing Gifts v1.1: How the computer world got this way, by Ted Nelson, ISBN: 978-0-578-00438-9, Published by Mindful Press, 2009, distributed through Lulu.Com, <http://www.lulu.com>

In a short but interesting book, computer pioneer Ted Nelson takes a very broad look at the origins and evolution of many of the basic ideas that underpin today's computer industry. The emphasis is on concepts and technologies rather than the success of individuals, the companies they founded, and the shape of the computer industry. This approach differentiates the book from other accounts, such as Robert X. Cringley's *Accidental Empires* and Martin Campbell-Kelly's *From Airline Reservations to Sonic the Hedgehog*.

Although the book is suitable for a fairly broad readership, an appreciation of the current makeup of the industry is helpful in understanding the significance of some of Nelson's ideas.

Organization

Geeks Bearing Gifts is divided into 60 short chapters, arranged in chronological order from the time the ideas originated, rather than when they appeared in fully developed form (indeed many are still developing). In the initial chapters Nelson covers topics such as language, alphabets, and encryption before moving on to examine the origins of computing. He then examines the contribution of pioneers from both inside and outside the United States, giving more credibility to contributors from outside of the United States than is normal.

As would be expected, Nelson deals in some detail with the topic of information presentation, in particular the origins of hypertext and associated developments such as *Xanadu* and the World Wide Web. He discusses the differences between these technologies, spending some time reflecting on his attempts to develop *Xanadu* at Brown University; he suggests that many of the deficiencies of the Web come from misdirection of that phase of the project.

Nelson next examines a wide selection of topics ranging from networks (both local and the Internet), object-orientated programming, and early desktop machines, before reaching the pivot point of his book: the UNIX operating system. He chose UNIX as the fulcrum of his analysis because he believes "so much led into it and so much has resulted from it."

Nelson next considers PUI (the PARC user interface), PCs, the role of the Microsoft and Apple operating systems and their evolution, the influence of the spreadsheet, the Internet, browsers, the Internet crash, and the current major companies in computing. He explores the promise, hype, and reality of the Web 2.0 model and its likely influence. (PARC stands for the Xerox *Palo Alto Research Center*.)

The last two chapters are summaries and thought guides. The first of these suggests that it is people and ideas rather than technology that advance the computer industry and that the myth of technological necessity has stifled imagination. The final chapter illustrates what the book is about—the disagreements and decisions that have made the technical world what it is today.

Synopsis

Nelson captures most of the important developments in the computer industry, although he acknowledges that in 199 pages it is possible to tell the reader only a little of where the software ideas come from and what they are. He sets out to show how varied and conflicting the initiatives that have propelled the evolution of computer technology have been, exposing the “ideas, disagreements, manoeuvres, forgotten possibilities, and politics.”

The book reads like a collection of themed essays, rather than a coherent sequence of stories. Nonetheless it is both informative and thought-provoking.

The Author

Ted Nelson is considered to be a radical thinker; he is one of the pioneers of the computer industry initiating the Xanadu project, which was started in the early 1960s with the objective of developing a computer network with a simple user interface. He is credited with inventing the term “hypertext.”

He holds a first degree in philosophy, a Masters in sociology, and a Doctorate in Media and Governance. Among his honors are a visiting fellowship at the Oxford Internet Institute and a Fellowship of Wadham College, Oxford; in addition, France has knighted him as “Officier des Arts et Lettres.” Visit:

http://en.wikipedia.org/wiki/Ted_Nelson

and

<http://www.ibiblio.org/pioneers/nelson.html>

...for more information.

—Edward Smith, BT, UK

edward.a.smith@btinternet.com

Read Any Good Books Lately?

Then why not share your thoughts with the readers of IPJ? We accept reviews of new titles, as well as some of the “networking classics.” In some cases, we may be able to get a publisher to send you a book for review if you don’t have access to it. Contact us at ipj@cisco.com for more information.

Fragments

RIPE Announces IPv6 Website

The RIPE NCC recently announced the launch of the *IPv6 Act Now!* website. Available at www.IPv6ActNow.org, the website explains IPv6 in terms that everyone can understand and provides a variety of useful information aimed at promoting the global adoption of IPv6. The site is designed for anyone with an interest in IPv6, including network engineers, company directors, law enforcement agencies, government representatives and civil society. The content is regularly updated and includes:

- Education, advice and opinions from the experts
- Latest IPv6-related news stories
- Videos and articles from Internet community leaders
- Current IPv4 exhaustion and IPv6 uptake statistics
- The RIPE community's statement on IPv6 deployment
- Information on community-developed IPv6 distribution policies
- Useful links to other sources of information about IPv6
- A forum for everyone to share experiences, ask questions and find answers

The site also includes contributions from other *Regional Internet Registries* (RIRs) and industry partners. If you have and comments or suggestions about the website, please contact:

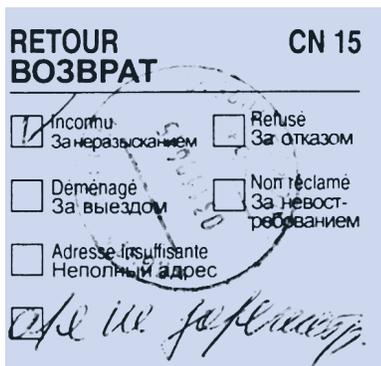
ipv6actnow@ripe.net

Four-byte AS numbers from APNIC

From July 1, 2009, the *Asia Pacific Network Information Centre* (APNIC) will assign four-byte *Autonomous System* (AS) numbers by default when receiving requests. Two-byte AS numbers will only be assigned if the applicant can demonstrate that a four-byte only AS number is unsuitable. This change marks the next phase of the transition to four-byte AS numbers. The final phase begins in January 2010, when APNIC will cease to make any distinction between two-byte and four-byte AS numbers, and will operate AS number assignments from an undifferentiated four-byte AS number pool. For more information please see: <http://icons.apnic.net/asn>

Please Tell Us When You Move

We receive large quantities of undeliverable copies of *The Internet Protocol Journal*. For international mailings, the returned mail piece usually includes a standard CN 15 label, an example of which is shown here. We have an extensive collection of CN 15 labels from all over the world, but we would much rather ensure that your journal is delivered to the correct address. So, if you're moving your home or office, please use the online subscription system to update your details, or just send an e-mail message to ipj@cisco.com with the new information. You can also suspend paper delivery and read IPJ online if you wish.



Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles (“What is...?”), as well as implementation/operation articles (“How to...”). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, “modem tax,” and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ contains standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US\$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at ole@cisco.com

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.



The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD
U.S. Postage
PAID
PERMIT No. 5187
SAN JOSE, CA

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
Distinguished Career Professor of Computer Science and Public Policy
Carnegie Mellon University, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, General Chair Person, WIDE Project
Vice-President, Keio University
Professor, Faculty of Environmental Information
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

*The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc. www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com*

Copyright © 2009 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners.

Printed in the USA on recycled paper.

