

Overview of 2000 IAB Wireless Internetworking Workshop

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document provides an overview of a workshop held by the Internet Architecture Board (IAB) on wireless internetworking. The workshop was hosted by Nokia in Mountain View, CA, USA on February 29 thru March 2, 2000. The goal of the workshop was to assess current and future uses of Internet technology in wireless environments, to make recommendations on research and standardization tasks to improve acceptance of Internet network and transport protocols in wireless environments, and to evaluate methods to improve communication and collaboration among Internet standards working groups and those of the telephony and wireless sectors. This report summarizes the conclusions and recommendations of the IAB on behalf of the IETF community.

Comments should be submitted to the IAB-Wireless-Workshop@ietf.org mailing list.

Table of Contents

1	Introduction	3
2	Presentation Overview	4
3	Discussion and Observations	9
3.1	Discussion on "Walled Garden" Service Model	9
3.2	Discussion on Mobility and Roaming	10
3.2.1	Discussion on Mobility and Roaming Model	11
3.2.2	Discussion on Mobility and Roaming Protocols	11
3.2.3	Discussion on Mobility and Roaming Services	12
3.3	Discussion on Security Model	12
3.3.1	Discussion on User Identity	12
3.3.2	Discussion on WAP Security	13

3.3.3	Discussion on 3G Network Security	13
3.4	Discussion on Transports	14
3.4.1	Discussion on Link Characteristics and Mobility Effect on Transport	14
3.4.2	Discussion on WAP Transport	16
3.4.3	Discussion on IETF Transport Activities	16
3.5	Discussion on Aeronautical Telecommunication Network (ATN) Routing Policy.	17
3.6	Discussion on QoS Services	18
3.6.1	Discussion on "Last Leg" QoS	18
3.6.2	Discussion on Path QoS Discovery	19
3.7	Discussion on Header Compression	20
3.8	Discussion on Applications Protocols	21
3.9	Discussion on Proxy Agents	22
3.10	Discussion on Adoption of IPv6	22
3.11	Discussion on Signaling	23
3.12	Discussion on Interactions Between IETF and Other Standards Organizations	24
4	Recommendations	25
4.1	Recommendations on Fostering Interaction with Non- Internet Standards Organizations	25
4.2	Recommendations for Dealing with "Walled Garden" Model	26
4.3	Recommendations on IPv4 and IPv6 Scaling	27
4.4	Recommendations on IPv4 and IPv6 Mobility	28
4.5	Recommendations on TCP and Transport Protocols	29
4.6	Recommendations on Routing	31
4.7	Recommendations on Mobile Host QoS Support	32
4.8	Recommendations on Application Mobility	33
4.9	Recommendations on TCP/IP Performance Characterization in WAP-like Environment	33
4.10	Recommendations on Protocol Encoding	33
4.11	Recommendations on Inter-Domain AAA Services	34
4.12	Recommendations on Bluetooth	34
4.13	Recommendations on Proxy Architecture	34
4.14	Recommendations on Justifying IPv6-based Solutions for Mobile / Wireless Internet	35
5	Security Considerations	35
6	Acknowledgments	35
7	Bibliography	36
A	Participants	41
B	Author's Address	41
	Full Copyright Statement	42

1 Introduction

Wireless technology, including wireless LANs, data transfer over cellular radio (GSM, 3GPP, etc), and mobile operations from aircraft and near earth spacecraft are becoming increasingly important. Some market projections suggest that a mobile Internet in parallel with or augmenting the wired Internet may be comparable in size to the wired Internet as early as 2003.

The wireless operators have not, however, chosen to use IPv4, TCP, full HTTP/HTML, and other applications for a variety of reasons. These relate to edge device cost, bandwidth limitations, perceived protocol imperfections, unnecessary complexities, the chattiness of the application protocols, and network layer addressing issues. Unfortunately, this creates some serious issues at the wired/wireless demarcation: end to end operation is sacrificed, security is compromised, and automated content modification in some form becomes necessary. The IAB considers these to be serious fundamental issues, which will in time be a serious impediment to the usability of the combined Internet if not addressed.

The Internet Architecture Board (IAB), on February 29 thru March 2, 2000, held an invitational workshop on wireless internetworking. The goal of the workshop was to assess current and future uses of Internet technology in wireless environments, to make recommendations on research and standardization tasks to improve acceptance of Internet network and transport protocols in wireless environments, and to evaluate methods to improve communication and collaboration among Internet standards working groups and those of the telephony and wireless sectors.

The following topics were defined for discussion:

- + Local area wireless technologies
- + Cellular wireless technologies
- + Wireless Application Protocol (WAP)
- + Near-space and aviation wireless applications
- + Voice over IP (VoIP) over wireless networks
- + Security over wireless networks
- + Transport and QoS over wireless networks
- + Use of WWW protocols over wireless and small screen devices

- + Addressing requirements for wireless devices
- + Compression and bit error requirements for wireless networks

The fundamental question addressed in these discussion is "what are the issues, and what really needs to be done to unify the Internet below the application layer." Applications will also need to be addressed, but were perceived to be more than could be usefully discussed in a three-day workshop, and probably require different expertise.

Section 2 presents a concise overview of the individual presentations made during the workshop. References to more extensive materials are provided. Details on major discussion topics are provided in section 3. Section 4 presents the recommendations made to wireless operators, IRTF, and IETF on the architectural roadmap for the next few years. It should be noted that not all participants agreed with all of the statements, and it was not clear whether anyone agreed with all of them. However, the recommendations made are based on strong consensus among the participants. Finally, section 5 highlights references to security considerations discussed, appendix A lists contact information of workshop participants, and appendix B lists the author contact information.

2 Presentation Overview

Title: Overview of Wireless IP Devices (Network Implications...)

Presenter: Heikki Hammainen

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/hh-IABpub.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/hh-IABpub.ppt>

Overview:

Title: Overview of IEEE 802.11 Wireless LAN's & Issues Running IP over IEEE 802.11?

Presenter: Juha Ala-Laurila

Reference:

http://www.iab.org/IAB-wireless-workshop/talks/IEEE80211_IP.ppt

Overview:

Title: Overview of Bluetooth Wireless & Issues Running IP over Bluetooth?

Presenter: Pravin Bhagwat

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/BT-overview.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/BT-overview.ppt>

Overview:

Title: Overview of Cellular Data Systems & Approaches to more IP centric Cellular Data System

Presenter: Jonne Soinien

Reference:

http://www.iab.org/IAB-wireless-workshop/talks/Cellular_JSo.PDF,
http://www.iab.org/IAB-wireless-workshop/talks/Cellular_JSo.ppt

Overview:

Title: IP Packet Data Service over IS-95 CDMA

Presenter: Phil Karn

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/karn/index.htm>

Overview:

Title: Wireless Internet Networking

Presenter: Chih-Lin I

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/IAB000229.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/IAB000229.ppt>

Overview:

Title: Mobile IP in Cellular Data Systems

Presenter: Charlie Perkins

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/WLIP99.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/WLIP99.ppt>

Overview:

Title: Overview of WAP

Presenter: Alastair Angwin

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/iab-wap-1.pdf>

Overview:

Title: Mobile Wireless Internet Forum (MWIF)

Presenter: Alastair Angwin

Reference:

http://www.iab.org/IAB-wireless-workshop/talks/MWIF_TC_Presentation.PDF,
http://www.iab.org/IAB-wireless-workshop/talks/MWIF_TC_Presentation.ppt

Overview:

Title: Some WAP History

Presenter: Jerry Lahti

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/waphist.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/waphist.ppt>

Overview:

Title: Near-space Wireless Applications

Presenter: Mark Allman

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/allman-iab-wireless.pdf>,
<http://www.iab.org/IAB-wireless-workshop/talks/allman-iab-wireless.ps>

Overview:

Title: Air Traffic / Aviation Wireless

Presenter: Chris Wargo

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/wargo-talk.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/wargo-talk.ppt>

Overview:

Title: VoIP over Wireless

Presenter: Christian Huitema

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/iab-wless-voip.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/iab-wless-voip.ppt>

Overview:

Title: Security Issues in Wireless Networks and Mobile Computing

Presenter: N. Asokan

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/mobile-security.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/mobile-security.ppt>

Overview:

Title: Security for Mobile IP in 3G Networks

Presenter: Pat Calhoun

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/mip-sec-3g.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/mip-sec-3g.ppt>

Overview:

Title: On Inter-layer Assumptions (A View from the Transport Area)

Presenter: Mark Handley

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/handley-wireless.pdf>,
<http://www.iab.org/IAB-wireless-workshop/talks/handley-wireless.ps>

Overview:

Title: Does current Internet Transport work over Wireless?

Presenter: Sally Floyd

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/IAB-wireless-Mar00.pdf>,
<http://www.iab.org/IAB-wireless-workshop/talks/IAB-wireless-Mar00.ps>

Overview:

Title: QOS for Wireless (DiffServ, IntServ, other?)

Presenter: Lixia Zhang

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/zhang-feb-IAB.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/zhang-feb-IAB.ppt>

Overview:

Title: Do current WWW Protocols work over Wireless and Small Screen Devices?

Presenter: Gabriel Montenegro

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/wireless-www.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/wireless-www.ppt>

Overview:

Title: Compression & Bit Error Requirements for Wireless

Presenter: Mikael Degermark

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/iab-hc.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/iab-hc.ppt>

Overview:

Title: Addressing Requirements for Wireless Devices & IPv6

Presenter: Bob Hinden

Reference:

<http://www.iab.org/IAB-wireless-workshop/talks/Addressing-IPv6.PDF>,
<http://www.iab.org/IAB-wireless-workshop/talks/Addressing-IPv6.ppt>

Overview:

3 Discussion and Observations

During the workshop presentations a number of issues were discussed and observations made. The following sections 3.1 -- 3.12 summarize these discussion and observations. Rather than organizing the material linearly by presentation, it is grouped according to common "themes" and issues.

3.1 Discussion on "Walled Garden" Service Model

Presentations from members involved in the cellular wireless (3GPP, 3G.IP, MWIF) and WAP environments quickly illustrated a significant difference in protocol specification and service models from that typically assumed by the Internet community. These communities focus on defining a profile (set of protocols and operational parameters) that combine to provide a well defined user service. In addition, the carriers typically prefer to have complete (or as much as possible) control over the entire service, including user access device, transmission facilities, and service "content". This style of service model appears to have been inherited from the classic telephony provider model. The term "walled garden" was coined to describe the resulting captive customer economic and service model. That is, the user is constrained within the limits of the service provided by the carrier with limited ability to extend features or access services outside the provider. The "walled garden" service model is in stark contrast to the "open" service assumed in the Internet. The application, access device, and service content may each be controlled by a different entity, and the service provider is typically viewed as little more than a "bit pipe".

Additionally, specification typically define a standalone protocol or application rather than the set of features and interoperation with other components required to deploy a commercial service.

Some discussion focused on whether cellular carriers could be persuaded to transition toward the Internet "open" service model. Responses indicated that there was little hope of this as carriers will always fight being reduced to a "bit pipe", fearing they cannot sustain sufficient revenues without the value added services. An additional point raised was that the closed model of the "walled garden" simplifies a number of issues, such as security, authorization, and billing when the entire network is considered secured and controlled under a single administration. These simplification can eliminate roadblocks to service deployment before scalable, interdomain solutions are available.

Even though there seems little hope of evolving carriers away from the "walled garden" service in the short term, there was significant value in recognizing its presence. This led to observations that "walled garden" Internet-based services will operate somewhat like current intranet services. Also, mechanisms should be investigated to simplify interoperation and controlled access to the Internet. Finally, the difference between Internet protocol specification contrasted to service profiles highlights some of the confusion those in the telephony environment encounter when attempting to incorporate Internet capabilities.

Much of the current work in extending Internet-based services to cellular customers has focused on data services such as email or web access. One observation on the reluctance of carriers to release any control over services was that this may be an impediment to adoption of Internet-based voice services. Current work on voice over IP (VoIP) and call signaling (SIP [30]) loosens control over these services, much of the functionality is moved into the SIP agent with the carrier being reduced to an access provider (i.e., "bit pipe").

3.2 Discussion on Mobility and Roaming

An inherent characteristic of wireless systems is their potential for accommodating device roaming and mobility. Some discussion focused on the model of mobility presented to the user. There was also considerable interest and discussion on protocols employed, using cellular telephony and/or IP-based solutions. Finally, there was some interest in exploring new services enabled by mobility.

3.2.1 Discussion on Mobility and Roaming Model

There was considerable discussion and concern over what style of mobility and roaming needs to be supported. Current usage in the Internet is dominated by the mode where a user performs some actions at one location, then shuts down and moves, followed by restart at a new location.

3G.IP uses the term "macro mobility" to describe this mode.

The discussion attempted to discern whether the current mode of usage is a perceived limitation introduced by current protocols. A clear consensus could not be achieved. There was agreement that introduction of this "macro mobility" roaming is a worthwhile first step. However, that was immediately followed by questions on whether it is a sufficient first step, and warning not to stop at this level. There seems significant issues for continued investigation related to enabling continual usage of a device during roaming ("micro mobility") and the ability to retrieve previous connections after a roaming event.

3.2.2 Discussion on Mobility and Roaming Protocols

Selection between cellular and IP protocols in support of roaming provided another topic for significant discussion. Cellular operators have already deployed protocols providing significant support for roaming. This has led several efforts, such as 3GPP and 3G.IP, toward architecture relying on telephone system for all mobility support, hiding roaming from the IP layer.

Arguments for cellular-based roaming centered on concerns about the mobile IP model. There was concern that home agent and foreign agent involvement in delivery might introduce bottleneck, and the perception that mobile IP handoff is too slow. A rebuttal offered was that IETF mobileip working group is introducing hierarchy and route optimization to improve performance and robustness [50], and there was disagreement on the point regarding slow handoff under mobile IP.

Detriments to the cellular-based roaming include the lack of IP support out to the mobile device and the added tunneling protocols and overhead required. Additionally, roaming is less well defined when traversing service provider boundaries and may involve highly non-optimal forwarding path. There appears significant work remaining to reach convergence on opinions, and additional investigation to support roaming across cellular, WLAN, and IP boundaries.

3.2.3 Discussion on Mobility and Roaming Services

3G.IP mobility model is primarily focused on providing ubiquitous service across a range of access media. However, the presentation also highlighted a desire to develop new "location based" services. Examples presented include locating nearby services or receiving advertisement and solicitations from nearby business.

There are several Internet protocols defined, such as anycast service [47] and SLP [28], that may aid in developing location based services. However, there was considerable frustration on the part of 3G.IP in that there appears little commercial support of these protocols, and even less direction on how to assemble and coordinate the required protocols to deploy the desired services.

This exchange illustrated the disconnect between interpreting Internet standards and telephony service profiles. First, in the Internet many protocols are defined but many are optional. Protocol support is typically driven by market demand, which can lead to "chicken and egg" problem. Secondly, individual protocols and applications are developed rather than complete profile to compose a commercial service. For this service, evaluating the usage and scalability of service discovery protocols appears to be an area open for further investigation.

3.3 Discussion on Security Model

Mobility and wireless environments introduce many complexities and potential attacks to user authentication and privacy. In addition to the discussion presented below, there was an overriding statement made regarding the methodology that must be followed for all security protocol development. It was felt quite strongly that the only chance for success is that the definition be done in a public forum, allowing full disclosure of all algorithms and thorough review by security experts. Stated an alternate way, defining protocols in a closed forum relying on cellphone manufacturers, or other non-experts on IP security, is very likely to create security exposures.

3.3.1 Discussion on User Identity

Storage of user identity can have significant effect on device usage and device portability. Discussion focused on whether identity should be tied to the mobile device or a transferable SIM card. Fixing identification with the device may simplify manufacture and provide some tamper resistance, however it makes it very difficult to deploy a public device taking on the identity of the user. These alternative also affect transfer of identity and configuration state on device replacement or upgrade.

A related topic revolves around the user desire to employ a single device but to take on a different identity and privilege based on the usage at hand (e.g., to gain corporate access, home access, or Internet access). The ability and ease of assuming these multiple identities may be highly dependent on the model of identity integration, as discussed above. Discussion highlighted potential pitfalls based on tying of device and user identities. IPsec use of device IP address inhibits roaming capabilities as the address may change based on location, and precludes distinguishing identity and capabilities for current usage. IPsec requires additional work to accommodate this added flexibility.

A final topic of discussion on user identity establishment was whether possession of the device is sufficient, or whether the user should be required to authenticate to the device. In the real world the first alternative is exemplified by the credit card model, while the second is more analogous to the ATM card where the user must also provide a PIN code. Both models seem useful in the real world, and it's likely both will have uses in wireless networking.

3.3.2 Discussion on WAP Security

WAP wireless transport security (WTLS) is based on TLS [20], with optimized handshake to allow frequent key exchange. The security service employs a "vertical" integration model, with protocol components throughout the network stack. Some argued that this is the wrong model. A better approach may have been a security layer with well defined interfaces. This could allow for later tradeoffs among different protocols, driven by market, applications, and device capabilities.

Additional statements argued that the WAP security model illustrates dangers from optimizing for a limited usage domain ("walled garden"). Content provider systems requiring security (e.g., banks) must deploy a special WAP proxy, which breaks the model of a single WAP "domain". Similar issues are inherent in gatewaying to the Internet.

3.3.3 Discussion on 3G Network Security

The existing GSM/GPRS model uses long term shared secrets (embedded in SIM card) with one-way authentication to the network, and with privacy only provided on the access link. This is an example where the "walled garden" service model has an advantage. Complete control over the service access devices and network greatly reduces the range of security concerns and potential attacks.

Future 3GPP and 3GPP2 plan to push IP all the way out to the wireless device. An observation is that this results in more potential for exposure of signaling and control plane to attacks. Desire is to perform mutual authentication and securing of the network. This is a difficult problem with additional issues remaining to be solved; however the statement was made that relying on IP and open standards is more likely to produce a provably secure network than former reliance on SS7 protocols and obscurity.

Completing support for the security requirements of the 3GPP/3GPP2 network seems to require resolving issues in two primary areas, AAA services and mobile IP. AAA is required for authentication, authorization, and billing. Remaining issues center around cross domain AAA, authentication using PKI, and there was considerable aversion to use of IPsec and IKE protocols due to perceived overhead and delay. Mobile IP issues revolve around solutions to reduce the security associations required between mobile node and home agent, mobile node and foreign agent, and the home and foreign agent. An interim solution being investigated involves use of a RADIUS server [56]; however, there are concerns with repeated dynamic key generation on each handoff or hiding some details of handoffs, which may violate assumptions in mobile IP protocol [48]. Evaluating requirements and addressing all of these open issues appears to be an excellent opportunity for mutual cooperation on open standardization and review.

3.4 Discussion on Transports

Discussion on transport protocols touched on a broad range of issues. Concerns ranged from the effects of wireless link characteristics and mobility effect on TCP, to development of new transport protocols such as WAP Wireless Transaction Protocol (WTP). In addition, a significant amount of time was spent reviewing ongoing efforts within the IETF on TCP transport enhancements and investigation of new transports.

3.4.1 Discussion on Link Characteristics and Mobility Effect on Transport

TCP makes assumptions on loss as congestion indication. The statement was made that TCP was designed for links with about 1% corruption loss, and provided that constraint is met then TCP should function properly. Presentation on IS-95 CDMA-based data service showed that it conditions line to provide 1--2% error rate with little correlation between loss. Similar conditioning and Forward Error Correction (FEC) mechanisms may be appropriate for other wireless and satellite systems [4]. This may not be true for all wireless media, but it was interesting in the fact that it indicates

TCP should work properly on many wireless media. However, the amount of discussion and suggestions on TCP performance optimizations showed that there can be a considerable gap between merely working and working well.

One issue raised several times was related to the effects of non-congestive loss on TCP performance. In the wireless environment non-congestive loss may be more prevalent due to corruption loss (especially if the wireless link cannot be conditioned to properly control error rate) or an effect of mobility (e.g., temporary outage while roaming through an area of poor coverage). These losses can have great detrimental effect on TCP performance, reducing the transmission window and halving the congestion window size. Much of the discussion focused on proposing mechanisms to explicitly indicate a non-congestive loss to the TCP source. Suggestions included a Non-Congestive Loss Indication (NCLI) sent for instance when packet corruption loss is detected, or sending a Source Encourage (SE) to stimulate source transmission at the end of an outage. In addition to data corruption, wireless links can also experience dropouts. In this situation any active TCP sessions will commence periodic retransmissions, using an exponentially increasing back-off timer between each attempt. When the link becomes available it may be many seconds before the TCP sessions resume transmission. Mechanisms to alleviate this problem, including packet caching and triggered retransmission were discussed. The more generic form of all of these mechanisms is one that allows the state of the layer two (datalink) system to signal to the TCP session its current operating mode. Developing a robust form of such a signaling mechanism, and integrating these signals into the end-to-end TCP control loop may present opportunities to improve TCP transport efficiency for wireless environments.

TCP improvements have been incorporated to support "long" links (i.e., those with large delay and bandwidth characteristics) [36], however considerable expertise may still be required to tune socket buffers for maximum performance. Some work has been done on auto-tuning buffers, which shows promise [58]. An additional problem with large windows and auto-tuning is the added header overheads. This may exasperate the problems of running TCP over low bandwidth links. Suggestions included to explore dynamic negotiation of large window extensions in the middle of a connection to alleviate these issues. A final issue raised with regardport (see discussion below in section 3.4.3).

There was also concern regarding mobility effects on TCP performance. TCP has implicit assumptions on bounding propagation delay. If delay exceeds the smoothed round trip time plus four times the round trip variance then the segment is considered lost, triggering the normal

backoff procedures. Could these assumptions be violated by segment loss or duplication during handoff? Work on D-SACK [25] may alleviate these worries, detecting reordering and allowing for adaptive DUP-ACK threshold. Finally, there was suggestion it might be appropriate to adapt (i.e., trigger slow start) immediately after mobile handoff on the assumption that path characteristics may differ.

3.4.2. Discussion on WAP Transport

WAPF considered TCP connection setup and teardown too expensive in terms of bit overhead and latency when required for every transaction. WAPF developed the Wireless Transaction Protocol (WTP), with some inspiration from T/TCP [12]. WTP offers several classes of service ranging from unconfirmed request to single request with single reply transaction. Data is carried in the first packet and 3-way handshake eliminated to reduce latencies. In addition acknowledgments, retransmission, and flow control are provided.

Discussion on WTP centered on assessing details on its operation. Although it incorporates mechanisms for reliability and flow control there was concern that it may miss critical or subtle transport issues learned through years of Internet research and deployment experience. One potential area for disaster appeared to be the use of fixed retransmission timers and lack of congestion control. This gave rise to suggestions that the IETF write up more details on the history and tradeoffs in transport design to aid others doing transport design work, and secondly that the IETF advocate that the congestion control is not optional when using rate adaptive transport protocols.

The remaining discussion on WAP transport primarily focused on ways to share information. It was suggested that any result from WAPF study of TCP shortcomings that led to its rejection might be useful for IETF review as inputs for TCP modifications. Similar comments were raised on study of T/TCP shortcomings and its potential exposure to Denial of Service (DoS) attacks. It was also encouraged that the WAPF members participate in the IETF directly contribute requirements and remain abreast of current efforts on evolving TCP operation and introduction of new transport (see discussion below in section 3.4.3.).

3.4.3 Discussion on IETF Transport Activities

Discussion on transport work in the IETF presented a large array of activities. Recent work on transport improvement includes path MTU, Forward Error Correction (FEC), large windows, SACK, NewReno Fast Recovery, ACK congestion control, segment byte counting, Explicit Congestion Notification (ECN), larger initial transmit windows, and

sharing of related TCP connection state [3,4,5,6,24,25,43,53,63]. Work on new transports includes SCTP [61] in the IETF Signaling Transport (sigtran) working group and TCP-Friendly Rate Control (TFRC) [1] by researchers at ACIRI. SCTP provides a reliable UDP-like protocol supporting persistent associations and in-order delivery with congestion control. TFRC is targeted at unreliable, unicast streaming media. Finally, work in the IETF End-point Congestion Management (ecm) working group is looking at standardizing congestion control algorithms, and work in the Performance Implications of Link Characteristics (pilc) working group is characterizing performance impacts of various link technologies and investigating performance improvements.

This vast array of ongoing research and standards development seemed a bit overwhelming, and there was considerable disagreement on the performance and applicability of several TCP extensions. However, this discussion did raise a couple of key points. First, transport work within the Internet community is not stagnant, there is a significant amount of interest and activity in improvement to existing protocols and exploration of new protocols. Secondly, the work with researchers in satellite networking has demonstrated the tremendous success possible in close collaboration. The satellite networking community was dissatisfied with initial TCP performance on long delay links. Through submission of requirements and collaborative investigation a broad range of improvements have been proposed and standardized to address unique characteristics of this environment. This should hopefully set a very positive precedent to encourage those in the wireless sector to pursue similar collaboration in adoption of Internet protocols to their environment.

3.5 Discussion on Aeronautical Telecommunication Network (ATN) Routing Policy

The Aeronautical Telecommunication Network (ATN) has goals to improve and standardize communications in the aviation industry. This ranges across air traffic management and control, navigation and surveillance, all the way up to passenger telephone service and entertainment. This also involves integration of both fixed ground segments and mobile aircraft. Supporting the ATN architecture using Internet protocols may introduce additional requirements on the routing infrastructure.

Current ATN views each aircraft as an autonomous network (AS) with changing point of attachment as it "roams" through different airspace. Addressing information associated with the aircraft is fixed, which makes route aggregation difficult since they're not related to topology, and also increases the frequency of updates. Additionally, the aircraft may be multiply attached (within coverage

of multiple ground and space-based access networks), requiring routing policy support for path selection. Finally, QoS path selection capabilities may be beneficial to arbitrate shared access or partition real-time control traffic from other data traffic.

Initial prototype of ATN capabilities have been based on ISO IDRP [33] path selection and QoS routing policy. There was some discussion whether IDRP could be adopted for use in an IP environment. There was quick agreement that the preferred solution within the IETF would be to advance BGP4++ [8,54] as an IDRP-like replacement. This transitioned discussion to evaluation of ATN use of IDRP features and their equivalent to support in BGP. Several issues with BGP were raised for further investigation. For example, whether BGP AS space is sufficient to accommodate each aircraft as an AS? Also issues with mobility support; can BGP provide for dynamically changing peering as point of attachment changes, and alternative path selection policies based on current peerings? A significant amount of additional investigation is required to fully assess ATN usage of IDRP features, especially in the QoS area. These could lead to additional BGP requirements, for instance to effect different prioritization or path selection for aircraft control vs. passenger entertainment traffic.

3.6 Discussion on QoS Services

Enabling support for voice and other realtime services along with data capabilities requires Quality of Service (QoS) features to arbitrate access to the limited transmission resources in wireless environment. The wireless and mobile environment requires QoS support for the last leg between the mobile device and network access point, accommodating roaming and unique characteristics of the wireless link.

In addition to the discussion presented below, it was felt quite strongly that it is critical any QoS facility be provided as an underlying service independent of payload type. That is, there should be no built in knowledge of voice or other application semantics. This results in a feature that can be leveraged and easily extended to support new applications.

3.6.1 Discussion on "Last Leg" QoS

Discussion on voice over IP (VoIP) emphasized that (wireless) access link is typically the most constrained resource, and while contention access (CSMA) provides good utilization for data it is not ideal for voice. Two models were identified as potential solution in VoIP architecture. The first is to have the wireless device directly signal the local access router. A second alternative is to have the

call control element (SIP agent [30]) "program" the edge router. This tradeoff seemed to be an area open for additional investigation, especially given the complications that may be introduced in the face of mobility and roaming handoffs. This appears a key component to solve for success in VoIP adoption.

Work within the IEEE 802.11 WLAN group identified similar requirements for QoS support. That group is investigating a model employing two transmission queues, one for realtime and one for best-effort traffic. Additional plans include mapping between IP DiffServ markings [14,46] and IEEE 802 priorities.

The statement was also made that QoS over the wireless link is not the fundamental problem, rather it is handling mobility aspects and seamless adaptation across handoffs without service disruption. There were concerns about mechanisms establishing per-flow state (RSVP [13]). Issues include scaling of state, and signaling overhead and setup delays on roaming events. DiffServ [9] approach allows allocating QoS for aggregate traffic class, which simplifies roaming. However, DiffServ requires measurement and allocation adjustment over time, and policing to limit amount of QoS traffic injected.

3.6.2 Discussion on Path QoS Discovery

The HDR high speed wireless packet data system under development at Qualcomm highlights unique characteristics of some wireless media. This system provides users a channel rate between 38.4Kb/s and 2.4Mb/s, with throughput dependent on channel loading and distance from network access point. This gave rise to considerable discussion on whether it might be possible to discover and provide feedback to the application regarding current link or path QoS being received. This might enable some form of application adaptation.

In the case of the HDR system it was indicated that no such feedback is currently available. Additionally, it was argued that this is in accord with the current Internet stack model, which does not provide any mechanisms to expose this type of information. Counter arguments stated that there are growing demands in Internet QoS working groups requesting exposure of this type of information via standardized APIs. Members working on GPRS protocols also indicated frustration in deploying QoS capabilities without exposure of this information. This clearly seemed a topic for further investigations.

A final area of discussion on QoS discovery focused on the question of how a server application might find out the capabilities of a receiver. This could allow for application adaptation to client device and path characteristics. One suggestion proposed use of RSVP payload, which is able to transport QoS information. A second

alternative is to push capability exchange and negotiation to the application layer. Discussion on this topic was brief, as application issues were deemed outside the workshop charter, however this also seems an area open for future investigation.

3.7 Discussion on Header Compression

A critical deterrent to Internet protocol adoption in the highly band-width constrained wireless cellular environment is the bit overhead of the protocol encoding. Examples presented highlighted how a voice application (layered over IP [52,19], UDP [51], and RTP [57]) requires a minimum of 40 bytes of headers for IPv4 or 60 bytes for IPv6 before any application payload (e.g., 24 byte voice sample). This overhead was also presented as a contributing factor for the creation of WAP Wireless Datagram Protocol (WDP) rather than IP for very low datarate bearers.

Discussion on header compression techniques to alleviate these concerns focused on work being performed within the IETF Robust Header Compression (rohc) working group. This working group has established goals for wireless environment, to conserve radio spectrum, to accommodate mobility, and to be robust to packet loss both before the point where compression is applied and between compressor and decompressor. Additional requirements established were that the technique be transparent, does not introduce additional errors, and that it is compatible with common protocol layerings (e.g., IPv4, IPv6, RTP/UDP/IP, TCP/IP).

The primary observation was that this problem is now largely solved! The working group is currently evaluating the ROCCO [38] and ACE [42] protocols, and expects to finalize its recommendations in the near future. It was reported that these encodings have a minimum header of 1 byte and result in average overhead of less than 2 bytes for an RTP/UDP/IP packet. There is some extra overhead required if transport checksum is required and some issues still to be analyzed related to interoperation with encryption and tunneling.

A detriment to IPv6 adoption often cited is its additional header overhead, primarily attributed to its larger address size. A secondary observation made was that it's believed that IPv6 accommodates greater header compression than IPv4. This was attributed to the elimination of the checksum and identification fields from the header.

Discussion on use of WWW protocols over wireless highlighted protocol encodings as another potential detriment to their adoption. A number of alternatives were mentioned for investigation, including use of a "deflate" Content-Encoding, using compression with TLS [20], or

Bellovin's TCP filters. Observation was made that it could be beneficial to investigate more compact alternative encoding of the WWW protocols.

3.8 Discussion on Applications Protocols

IETF protocol developments have traditionally taken the approach of preferring simple encode/decode and word alignment at the cost of some extra bit transmissions. It was stated that optimizing protocol encoding for bit savings often leads to shortcomings or limitations on protocol evolution. However, it was also argued that environments where physical limitations have an effect on transmission capacity and system performance may present exceptions where optimized encodings are beneficial. Cellular wireless and near-space satellite may fall into this category.

The WAP protocols exhibit several examples where existing Internet protocols were felt to be too inefficient for adoption with very low datarate bearer services and limited capability devices. The WAP Wireless Session Protocol (WSP) is based on HTTPv1.1 [23], however WSP incorporates several changes to address perceived inefficiencies. WSP uses a more compact binary header encoding and optimizations for efficient connection and capability negotiation. Similarly, the WAP Wireless Application Environment (WAE) uses tokenized WML and a tag-based browser environment for more efficient operation.

Additional requests for more efficient and compact protocol encodings, and especially improved capability negotiation were raised during discussion on usage of WWW protocols with wireless handheld devices.

Finally, work within the near-space satellite environment has pointed out other physical limitations that can affect performance. In this case the long propagation delays can make "chatty" protocols highly inefficient and unbearable for interactive use. This environment could benefit from protocols that support some form of "pipelining" operation.

There seemed broad agreement that many of these observations represent valid reasons to pursue optimization of protocol operations. Investigation of compact protocol encoding, capability negotiation, and minimizing or overlapping round trips to complete a transaction could all lead to improved application performance across a wide range of environments.

3.9 Discussion on Proxy Agents

Proxy agents are present in a number of the wireless and mobile architectures. They're often required to gateway between communication domains; terminate tunnel and translate between telephony system and Internet protocols (GPRS), or to escape the "walled garden" (WAP). In conjunction with limited capability handheld devices a proxy might be deployed to offload expensive processing such as public key operations, perform content filtering, or provide access to "backend" applications (e.g., email, calendar, database). In other cases the proxy may be required to work around protocol deployment limitations (e.g., NAT with limited IPv4 addresses).

The discussion on proxy agents primarily recognized that there are a range of proxy agent types. Proxies may operate by intercepting and interpreting protocol packets, or by hijacking or redirecting connections. Some types of proxy break the Internet end-to-end communication and security models. Other proxy architectures may limit system scalability due to state or performance constraints. There was some desire to conduct further study of proxy agent models to evaluate their effect on system operation.

3.10 Discussion on Adoption of IPv6

Projections were presented claiming 1200 million cellular (voice) subscribers, 600 million wired stations on the Internet, and over 600 million wireless data ("web handset") users by the year 2004. Right up front there was caution about these projections, especially the wireless data since it is highly speculative with little history. Secondly, there was some doubt regarding potential for significant revenues from user base over 1 billion subscribers; this may be pushing the limits of world population with sufficient disposable income to afford these devices. However, there was broad consensus that cellular and Internet services are going to continue rapid growth and that wireless data terminals have potential to form a significant component of the total Internet. These conclusions seemed to form the basis for many additional recommendations to push for adoption of IPv6 protocols in emerging (3G) markets.

In nearly all the presentations on 3G cellular network technologies discussion on scaling to support the projected large number of wireless data users resulted in strong advocacy by the Internet representatives for adoption of IPv6 protocols. There were some positive signs that groups have begun investigation into IPv6. For example, 3GPP has already defined IPv6 as an option in their 1998 and 1999 specifications (release R98 and R99), and are considering

specifying IPv6 as mandatory in the release 2000. The MWIF effort is also cognizant of IPv4 and IPv6 issues and is currently wrestling with their recommendations in this area.

Although there was limited positive signs on IPv6 awareness, indication is that there are long fights ahead to gain consensus for IPv6 adoption in any of the 3G standards efforts. There was considerable feedback that the telephony carriers perceive IPv6 as more difficult to deploy, results in higher infrastructure equipment expenses, and adds difficulty in interoperation and gatewaying to the current (IPv4) Internet. Arguments for sticking with IPv4 primarily came down to the abundance and lower pricing of IPv4-based products, and secondary argument of risk aversion; there is currently minimal IPv6 deployment or operational experience and expertise, and the carriers do not want to drive development of this expertise. Finally, some groups argue IPv4 is sufficient for "walled garden" use, using IPv4 private address space (i.e., the "net 10" solution).

One other area of concern regarding IPv6 usage is perceived memory and processing overhead and its effect on small, limited capability devices. This was primarily directed at IPv6 requirement for IPsec implementation to claim conformance. Arguments that continued increase in device capacity will obviate these concerns were rejected. It was stated that power constraints on these low-end devices will continue to force concerns on memory and processing overhead, and impact introduction of other features. There was no conclusion on whether IPsec could be made optional for these devices, or the effect if these devices were "non-compliant".

Emerging 3G cellular networks appear ideal environment for IPv6 introduction. IPv6 addresses scaling requirements of wireless data user projections and eliminates continued cobbling of systems employing (IPv4) private address space and NAT. This appears an area for IAB and Internet community to take a strong stance advocating adoption of IPv6 as the various 3G forums wrestle with their recommendations.

3.11 Discussion on Signaling

Discussion on signaling focused on call setup and control functions, and the effects of mobility. The 3G.IP group has investigated standardizing on either H.323 [32] or SIP [30]. Currently support seems to be split between the protocols, and neither seemed ideal without support for mobility. During discussion on VoIP it was presented that SIP does support mobility, with graceful handling of mobile handoff, updating location information with remote peer, and even simultaneous handoff of both endpoints. The problem with SIP adoption seems to be its slow standardization brought about by

focusing on the harder multicast model rather than expediting definition of a unicast "profile". There seems great need for IETF to expedite finalization of SIP, however some argued at this point it's likely many products will need to develop support for both SIP and H.323, and for their interoperation.

A short discussion was also raised on whether it is the correct model to incorporate the additional protocol mechanisms to accommodate mobility into the SIP signaling. An alternative model might be to build on top of the existing mobile IP handoff facilities. There was no conclusion reached, however it seemed an area for further investigation.

3.12 Discussion on Interactions Between IETF and Other Standards Organizations

There were many examples where non-IETF standards organizations would like to directly adopt IETF standards to enable Internet (or similar) services. For example IEEE 802.11 WLAN relies on adoption of IETF standards for mobile IP, end-to-end security, and AAA services. 3GPP is looking into the IETF work on header compression. WAPF derived its transport, security, and application environment from Internet protocols. At first glance these would seem successes for adoption of Internet technologies, however the decision to rely on IETF standards often introduced frustrations too.

One common theme for frustration is differences in standardization procedures. For instance, 3GPP follows a strict model of publishing recommendations yearly; any feature that cannot be finalized must be dropped. On the other hand the IETF working groups have much less formalized schedules, and in fact often seem to ignore published milestone dates. This has led to a common perception within other standards organizations that the IETF cannot deliver [on time].

A second area identified where IETF differs from other organizations is in publication of "system profile". For example defining interoperation of IPsec, QoS for VoIP and video conferencing, and billing as a "service". Wading through all the protocol specifications, deciding on optional features and piecing together the components to deliver a commercial quality service takes considerable expertise.

Thirdly, there was often confusion about how to get involved in IETF standards effort, submit requirements, and get delivery commitments. Many people seem unaware and surprised at how open and simple it is to join in IETF standardization via working group meetings and mailing list.

There wasn't really a large amount of discussions on ways to address these differences in standards practices. However, it did seem beneficial to understand these concerns and frustrations. It seemed clear there can be some benefits in improving communication with other standards organizations and encouraging their participation in IETF activities.

4 Recommendations

The IAB wireless workshop provided a forum for those in the Internet research community and in the wireless and telephony community to meet, exchange information, and discuss current activities on using Internet technology in wireless environments. However the primary goal from the perspective of the IAB was to reach some understanding on any problems, both technical or perceived deficiencies, deterring the adoption of Internet protocols in this arena. This section documents recommendations of the workshop on actions by the IAB and IESG, IRTF research efforts, and protocol development actions for the IETF to address these current deficiencies and foster wider acceptance of Internet technologies.

4.1 Recommendations on Fostering Interaction with Non-Internet Standards Organizations

A clear consensus of the workshop is that dialog needs to be improved. The Internet community should attempt to foster communication with other standards bodies, including WAPF, MWIF, 3GPP, 3G.IP, etc. The goal is to "understand each others problems", provide for requirements input, and greater visibility into the standardization process.

4.1.1

It was recommended to take a pragmatic approach rather than formalizing liaison agreements. The formalized liaison model is counter to the established Internet standards process, is difficult to manage, and has met with very limited success in previous trials. Instead, any relevant IETF working group should be strongly encouraged to consider and recommend potential liaison requirements within their charter.

4.1.2

It was recommended to avoid formation of jointly sponsored working groups and standards. Once again this has shown limited success in the past. The preferred mode of operation is to maintain separate standards organizations but to encourage attendance and participation of external experts within IETF proceedings and to avoid overlap.

An exception to this style of partitioning meeting sponsorship is less formal activities, such as BOFs. It was recommended that sponsoring joint BOF could be beneficial. These could enable assembly of experts from multiple domains early in the process of exploring new topics for future standards activities.

4.1.3

A principle goal of fostering communication with other standards organizations is mutual education. To help in achieving this goal recommendations were made related to documenting more of the history behind Internet standards and also in coordinating document reviews.

It was recommended that IETF standards groups be encouraged to create or more formally document the reasons behind algorithm selection and design choices. Currently much of the protocol design history is difficult to extract, in the form of working group mail archives or presentations. Creation of these documents could form the basis to educate newcomers into the "history" and wisdom behind the protocols.

It was recommended that mutual document reviews should be encouraged. This helps to disseminate information on current standards activities and provides an opportunity for external expert feedback. A critical hurdle that could severely limit the effectiveness of this type of activity is the intellectual property and distribution restrictions some groups place on their standards and working documents.

4.2 Recommendations for Dealing with "Walled Garden" Model

There are several perceived benefits to the "walled garden" (captive customer) model, similar to current deployment of "intranets". These range from simplified user security to "captive customer" economic models. There was disagreement on the extent this deployment model might be perpetuated in the future. However it is important to recognize this model exists and to make a conscious decision on how to accommodate it and how it will affect protocol design.

4.2.1

It was strongly recommended that independent of the ubiquity of the "walled garden" deployment scenario that protocols and architectural decisions should not target this model. To continue the success of Internet protocols at operating across a highly diverse and heterogeneous environment the IETF must continue to foster the adoption of an "open model". IETF protocol design must address seamless, secure, and scalable access.

4.2.2

Recognition that the "walled garden" model has some perceived benefits led to recommendations to better integrate it into the Internet architecture. These focused on service location and escape from the "walled garden".

It was recommended to investigate standard protocols for service and proxy discovery within the "walled garden" domain. There are already a number of candidate mechanisms, including static preconfiguration, DNS [22,27,44,45], BOOTP [18], DHCP [21], SLP [28], and others. Specific recommendations on use of these protocols in this environment can help foster common discovery methods across a range of access devices and ease configuration complexity.

It was recommended to investigate standard methods to transport through the garden wall (e.g., escape to the Internet). It seemed clear that a better model is required than trying to map all access over a HTTP [23] transport connection gateway. One suggestion was to propose use of IP!

4.3 Recommendations on IPv4 and IPv6 Scaling

Wireless operators are projecting supporting on the order of 10's to 100's million users on their Internet-based services. Supporting this magnitude of users could have severe scaling implications on use of the dwindling IPv4 address space.

4.3.1

There was clear consensus that any IPv4-based model relying on traditional stateless NAT technology [60] is to be strongly discouraged. NAT has several inherent faults, including breaking the Internet peer-to-peer communication model, breaking end-to-end security, and stifling deployment of new services [16,29,31]. In addition, the state and performance implications of supporting 10's to 100's million users is cost and technologically prohibitive.

4.3.2

Realm specific IP (RSIP) [10,11] has potential to restore the end-to-end communication model in the IPv4 Internet, broken by traditional NAT. However there was considerable reluctance to formally recommend this as the long term solution. Detriments to its adoption include that the protocol is still being researched and defined, and potential interactions with applications, QoS features, and security remain. In addition, added signaling, state, and tunneling has cost and may be technologically prohibitive scaling.

4.3.3

The clear consensus of the workshop was to recommend adoption of an IPv6-based solution to support these services requiring large scaling. Adoption of IPv6 will aid in restoring the Internet end-to-end communication model and eliminates some roaming issues. Adoption of IPv6 in this marketplace could also help spur development of IPv6 products and applications, and hasten transition of the Internet. It was recognized that some application gateways are required during transition of the IPv4 Internet, however it was felt that the scaling and roaming benefits outweighed these issues.

4.3.4

It was recommended that an effort be made to eliminate any requirement for NAT in an IPv6 Internet. The IAB believes that the IPv6 address space is large enough to preclude any requirement for private address allocation [55] or address translation due to address space shortage [15]. Therefore, accomplishing this should primarily require installing and enforcing proper address allocation policy on registry and service providers. It was recommended to establish policies requiring service providers to allocate a sufficient quantity of global addresses for a sites use. The feeling was that NAT should be easily eliminated provided efficient strategies are defined to address renumbering [17,62] and mobility [37] issues.

4.4 Recommendations on IPv4 and IPv6 Mobility

An inherent characteristic of wireless systems is their potential for accommodating device roaming and mobility. Scalable and efficient support of this mobility within Internet protocols can aid in pushing native IP services out to the mobile devices.

4.4.1

Several limitations were identified relating to current specification of mobile IPv4 [48]. Primary among these limitations is that mechanisms to support redundant home agents and failover are not currently defined. Redundant home agents are required to avoid single point of failure, which would require (proprietary) extensions. Additional deficiencies related to lack of route optimization, and tunneling and path MTU issues were also identified. Due to these limitations there was reluctance to recommend this as a solution.

4.4.2

It was recommended to encourage adoption of IPv6 mobility extensions [37] to support roaming capabilities in the wireless environment. IP mobility over IPv6 incorporates improvements to address several limitations of the IPv4-based mobility. The ability to use autoconfiguration for "care of" address improves robustness and efficiency. Additionally, path MTU is more easily adapted when a router forwards to a new "care of" address.

Building wireless roaming atop IPv6-based mobility may introduce IPv4/IPv6 transition issues unique to the mobile environment. It was recommended to add investigation of these issues to the charter of the existing IETF Next Generation Transition (ngtrans) working group, provided any mobile IP interoperation issues be identified.

4.4.3

Scalable and widespread authentication, authorization, and accounting (AAA) services are critical to the deployment of commercial services based on (wireless) mobile IP. Some work is progressing on definition of these standards for IP mobility [26,49]. However, due to the pivotal role of these protocols on the ability to deploy commercial services, it was recommended to make finalization of these AAA standards and investigation of AAA scalability as high priorities.

4.5 Recommendations on TCP and Transport Protocols

The wireless environment and applications place additional requirements on transport protocol. Unique link error and performance characteristics, and application sensitivity to connection setup and transaction semantics has led to "optimized" transports specific to each environment. These new transports often lack robustness found in Internet transport and place barriers to seamless gatewaying to the Internet. It was felt that better education on transport design and cooperation on Internet transport evolution could lead to significant improvements.

4.5.1

It was recommended that the IETF Transport Area (tsv) working group document why Internet transport protocols are the way they are. The focus should be on generic transport issues and mechanisms, rather than TCP specifics. This should capture usage and tradeoffs in design of specific transport mechanisms (e.g., connection

establishment, congestion control, loss recovery strategies, etc.), and document some of the history behind transport research in the Internet.

This "entry point" document into transport design is in direct support of the recommendations in section 4.1 to foster communication and mutual education. In addition it was deemed critical that the Internet community make it very clear that congestion control is not optional. Internet researchers have learned that optimizing for a single link or homogeneous environment does not scale. Early work by Jacobson [34,35], standardization of TCP congestion control [5], and continuing work within the IETF Endpoint Congestion Management (ecm) working group could provide excellent basis for education of wireless transport designers.

4.5.2

It was recommended that the IETF actively solicit input from external standards bodies on identifying explicit requirements and in assessing inefficiencies in existing transports in support of cellular and wireless environments. This has proven highly effective in identifying research topics and in guiding protocol evolution to address new operational environments, for instance in cooperation with groups doing satellite-based internetworking [4,6].

4.5.3

It was recommended that the IAB make wireless standards bodies aware of the existence, and get them active in, the IETF Transport Area (tsv) working group. This transport "catch all" could provide an excellent forum for workers outside the Internet community to propose ideas and requirements, and engage in dialog with IESG members prior to contributing any formal proposal into the IETF or incurring overhead of working group formation.

4.5.4

Mobile radio environments may often be subject to frequent temporary outages. For example, roaming through an area that is out of range of any base station, or disruptions due to base station handoffs. This violation of the congestive loss assumption of TCP can have severe detrimental effect on transport performance. It was recommended to investigate mechanisms for improving transport performance when these non-congestive loss can be detected. Areas for potential research identified include incorporation of "hints" to the sender providing Non-Congestive Loss Indication (NCLI) or stimulating transmission after link recovery via Source Encourage

(SE) message [39]. This likely falls to the auspice of the IETF Performance Implications of Link Characteristics (pilc) working group.

4.5.5

Many wireless applications require transaction semantics and are highly sensitive to connection establishment delays (e.g., WAP). However, it is still desirable to efficiently support streaming of large bulk transfers too. It was recommended to investigate tradeoffs in supporting these transaction and streaming connections. Potential areas for investigation include tradeoffs between minimal transaction transport and potential security and denial of service (DoS) attacks, mechanisms to piggyback data during connection establishment to eliminate round trip delays, or ways for endpoints to cooperate in eliminating setup handshake for simple transactions while providing switch-over to reliable streaming for bulk transfers.

4.5.6

It was recommended to look at (TCP) transport improvements specific to the wireless and mobile environment. An example is to investigate reattachable transport endpoints. This could allow for graceful recovery of a transport connection after a roaming or mobility event results in changes to one or both endpoint identifiers. Another area for potential investigation is to develop targeted uses of D-SACK [25]. D-SACK provides additional robustness to reordered packets, which may prove beneficial in wireless environment where packets are occasionally corrupted. Higher performance may be attainable by eliminating requirements on link-level retransmission maintaining in-order delivery within a flow.

4.6 Recommendations on Routing

Unique routing requirements may be introduced in support of wireless systems, especially when viewing the mobile component as an autonomous system (AS).

4.6.1

It was recommended that the IETF Routing Area commence investigation of extensions to the BGP protocol [54] to support additional policy features available within the ISO IDRP protocol [33]. The range of policy control desired includes adopting different identity or policies based on current point of attachment, and providing flexibility in path selection based on local policy and/or current

peer policy. These features could be used for instance in support of requirements established in the Aeronautical Telecommunication Network (ATN).

4.6.2

It was recommended that the IETF Routing Area commence investigation of extensions to the BGP protocol [54] to support additional QoS/TOS path selection features available within the ISO IDRP protocol [33]. The range of policies include differentiating service level or path selection based on traffic classes. An example, based on Aeronautical Telecommunication Network (ATN) requirements, might be differentiating path selection and service between airline control and passenger entertainment traffic.

4.7 Recommendations on Mobile Host QoS Support

Wireless link bandwidth is often scarce (e.g., cellular) and/or shared (e.g., IEEE 802.11 WLAN). Meeting application QoS needs requires accommodating these link characteristic, in addition to the roaming nature of mobile host. Specialized support may be required from the network layer to meet both link and end-to-end performance constraints.

4.7.1

It was recommended that the IETF Transport Area undertake investigation into providing QoS in the last leg of mobile systems. That is, between the mobile device and the network access point. This type of QoS support might be appropriate where the wireless link is the most constrained resource. A potential solution to investigate is to employ an explicit reservation mechanism between the mobile host and the access point (e.g., RSVP [13]), while relying on resource provisioning or more scalable DiffServ [9] technologies within the core.

4.7.2

It was recommended that the IETF Transport Area undertake investigation into end-to-end QoS when the path includes a mixture of wireless and wired technologies. This investigation could focus on mechanism to communicate QoS characteristics in cellular network to the core network to establish end-to-end QoS guarantees. An alternative investigation is to look into discovery problem of assessing current end-to-end performance characteristics, enabling for dynamic adaptation by mobile host.

4.8 Recommendations on Application Mobility

In a mobile environment with roaming, and mobile host disconnect and reconnect at different attachment point it may be desirable to recover an incomplete application session. It was recommended that the IRTF investigate application mobility at this level. The goal is to achieve a smooth recovery after a disconnect period; something more graceful than a "redial". Currently there does not appear to be sufficient information available within the network stack, this may require instantiation of some form of "session" layer.

4.9 Recommendations on TCP/IP Performance Characterization in WAP-like Environment

WAPF has gone to considerable effort to develop unique transport protocol and optimizations due to perception that TCP/IP protocol stack is too expensive. Much of this was predicated on WAP requirements to support very low datarate bearer services. It was recommended that members of the IRTF evaluate TCP/IP stack performance in WAP-like environment to quantify its behavior and applicability. The focus should include investigation of code and memory space requirements, as well as link usage to complete a single transaction for current WAP protocols and for both IPv4 and IPv6. This work should result in better characterization of TCP/IP performance in highly constrained devices and network, recommendations to the IETF on protocol enhancements to optimize performance in this environment, and recommendations to WAPF on suitability of deploying native IP protocols.

4.10 Recommendations on Protocol Encoding

IETF protocol developments have traditionally taken the approach of preferring simple encode/decode and word alignment at the cost of some extra bit transmissions. This overhead may prove too burdensome in some bandwidth constrained environments, such as cellular wireless and WAP. Work within the IETF Robust Header Compression (rohc) working group may go a long way to reducing some of these detriments to Internet protocols deployment. However, there may be potential for additional savings from investigation of alternative encoding of common Internet protocols. It was recommended that members of the IRTF evaluate general techniques that can be used to reduce protocol "verbiage". Examples might include payload compression techniques or tokenized protocol encoding.

4.11 Recommendations on Inter-Domain AAA Services

Commercial roaming and mobility services are likely to require exchange of authentication, authorization, and billing services spanning multiple domains (service providers). This introduces requirements related to establishing a web or hierarchy of trust across multiple autonomous domains. Standard protocols to specify and exchange usage policies and billing information must also be established. Some work is progressing on scoping out the issues and a framework [7,64]. However, there are significant issues to be solved to enable a scalable, Internet-wide solution. Due to the pivotal role of these protocols on the ability to deploy commercial services, it was recommended to make finalization of scalable inter-domain AAA as high priority within the IETF.

4.12 Recommendations on Bluetooth

Bluetooth protocols and devices were originally optimized for a narrow application space. However, there is interest in exploring the breadth to which protocol and device access can be extended. One particular area of interest is exploring integration into, or gatewaying access to, the Internet. It was recommended that the IETF pursue formation of a joint BOF to assemble experts from the IETF and Bluetooth communities to begin exploration of this problem. This is in direct support of the recommendations in section 4.1 to foster communication and mutual education.

4.13 Recommendations on Proxy Architecture

Proxy agents are often deployed to intercept and evaluate protocol requests (e.g., web cache, HTTP redirector, filtering firewall) or to gateway access between communication domains (e.g., traversing bastion host between private network and Internet or gatewaying between a cellular service and the Internet). There are a number of potential architectures when contemplating development and deployment of one of these proxy agent. It was recommended that members of the IRTF investigate taxonomy of proxy architectures and evaluate their characteristics and applicability. Each type of proxy should be characterized, for example, by its effect on Internet end-to-end model, and security, scaling, and performance implications. The results of this study can help educate developers and network operators on the range of proxy available and recommend solutions that are least disruptive to Internet protocols.

4.14 Recommendations on Justifying IPv6-based Solutions for Mobile / Wireless Internet

IPv6 was strongly recommended to address scaling (see section 4.3) and mobility (see section 4.4) issues in the future Internet dominated by large numbers of wireless and mobile devices. It was recommended that the IAB draft a formalized justification for these recommendations for adoption of IPv6-based solution. It was believed that the "The Case for IPv6" [40] document should form an excellent basis for this justification. In addition, documents highlighting architectural and operational pitfalls of continued reliance on IPv4 and NAT also provide excellent justification [29,31,59]. It was deemed urgent to submit these informational documents as inputs to other standards bodies (MWIF, 3GPP, 3G.IP), as many decisions are being made on Internet protocol adoption and this data could be highly influential.

5 Security Considerations

This workshop did not focus on security. However, mobility and wireless environment introduces additional complexities for security and potential attacks to user authentication and privacy. The presentations by Asokan and by Calhoun referenced in section 2 focused on security mechanisms in currently deployed cellular networks and evolution toward 3G cellular and IP networks. Discussion on the "walled garden" service model (see section 3.1) briefly mentions effects on simplifying security requirements. Section 3.3 raises a number of security issues related to wireless devices and mobility. These include alternatives for establishing user identity and capabilities, securing network infrastructure from attacks, and security associations required for mobile IP and AAA operation. Section 3.7 mentions interoperation issues between compression and encryption or tunneling, and finally section 3.9 highlight potential for proxy agent to be used to offload expensive crypto operations.

6 Acknowledgments

The author would like to thank all of the workshop participants for their feedback, encouragement, and patience during the writeup of this document. I would especially like to thank Brian Carpenter for prompt responses to questions on the document organization and content. Similarly, Charlie Perkins provided extensive feedback that dramatically improved and corrected statements throughout the report. Finally, Mikael Degermark, Sally Floyd, Heikki Hammainen, Geoff Huston, and Gabriel Montenegro contributed comments and responses to questions.

7 Bibliography

- [1] ACIRI. TCP-Friendly Rate Control. <http://www.aciri.org/tfrc>.
- [2] A. Aggarwal, S. Savage, and T. Anderson. Understanding the Performance of TCP Pacing. Proceedings of IEEE Infocom 2000, March 2000.
- [3] Allman, M., Floyd, S. and C. Partridge, "Increasing TCP's Initial Window", RFC 2414, September 1998.
- [4] Allman, M., Glover, D. and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", RFC 2488, January 1999.
- [5] Allman, M., Paxson, V. and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.
- [6] Allman, M., Dawkins, S., Glover, D., Griner, J., Tran, D., Henderson, T., Heidemann, J., Touch, J., Kruse, H., Ostermann, S., Scott, K. and J. Semke, "Ongoing TCP Research Related to Satellites", RFC 2760, February 2000.
- [7] Arkko, J., "Requirements for Internet-Scale Accounting Management", Work in Progress.
- [8] Bates, T., Chandra, R., Katz, D. and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 2283, February 1998.
- [9] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services" RFC 2475, December 1998.
- [10] Borella, M., et al., "Realm Specific IP: Framework", Work in Progress.
- [11] Borella, M., et al., "Realm Specific IP: Protocol Specification", Work in Progress.
- [12] Braden, R., "T/TCP -- TCP Extensions for Transactions Functional Specification", RFC 1644, July 1994.
- [13] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [14] Brim, S., Carpenter, B. and F. Le Faucheur, "Per Hop Behavior Identification Codes", RFC 2836, May 2000.

- [15] Carpenter, B., Crowcroft, J. and Y. Rekhter, "IPv4 Address Behaviour Today", RFC 2101, February 1997.
- [16] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [17] Crawford, M., "Router Renumbering for IPv6", RFC 2894, August 2000.
- [18] Croft, B. and J. Gilmore, "Bootstrap Protocol (BOOTP)", RFC 951, September 1985.
- [19] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [20] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.
- [21] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [22] Everhart, C., Mamakos, L., Ullman, R. and P. Mockapetris, "New DNS RR Definitions", RFC 1183, October 1990.
- [23] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [24] Floyd, S. and T. Henderson, "The NewReno Modification to TCP's Fast Recovery Algorithm", RFC 2582, April 1999.
- [25] Floyd, S., Mahdavi, J., Mathis, M. and M. Podolsky, "An Extension to the Selective Acknowledgment (SACK) Option for TCP", RFC 2883, July 2000.
- [26] Glass, S., Hiller, T., Jacobs, S. and C. Perkins, "Mobile IP Authentication, Authorization, and Accounting Requirements", RFC 2977, October 2000.
- [27] Gulbrandsen, A. and P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2052, October 1996.
- [28] Guttman, E., Perkins, C., Veizades, J. and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.
- [29] Hain, T., "Architectural Implications of NAT", RFC 2993, November 2000.

- [30] Handley, M., Schulzrinne, H., Schooler, E., and J. Rosenberg, "SIP: Session Initiation Protocol", RFC 2543, March 1999.
- [31] Holdrege, M. and P. Srisuresh, "Protocol Complications with the IP Network Address Translator (NAT)", Work in Progress.
- [32] International Telecommunication Union. Visual Telephone Systems and Equipment for Local Area Networks which provide a Non-guaranteed Quality of Service. Recommendation H.323, May 1996.
- [33] ISO/IEC. Protocol for Exchange of Inter-Domain Routeing Information among Intermediate Systems to support Forwarding of ISO 8473 PDUs. ISO/IEC IS10747, 1993.
- [34] V. Jacobson. Congestion Avoidance and Control. Computer Communication Review, vol. 18, no. 4 August 1988.
<ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z>.
- [35] V. Jacobson. Modified TCP Congestion Avoidance Algorithm. end2end-interest mailing list, April 30, 1990.
<ftp://ftp.isi.edu/end2end/end2end-interest-1990.mail>.
- [36] Jacobson, V., Braden, R. and D. Borman, "TCP Extensions for High Performance", RFC 1323, May 1992.
- [37] Johnson, D. and C. Perkins, "Mobility Support in IPv6", Work in Progress.
- [38] Jonsson, L., et al., "RObust Checksum-based header COmpression (ROCCO)", Work in Progress.
- [39] Karn, P., et al., "Advice for Internet Subnetwork Designers", Work in Progress.
- [40] King, S., et al., "The Case for IPv6", Work in Progress.
- [41] J. Kulik, R. Coulter, D. Rockwell, and C. Partridge. Paced TCP for High Delay-Bandwidth Networks. Proceedings of IEEE Globecom '99, December 1999.
- [42] Le, K., et al., "Adaptive Header ComprESSION (ACE) for Real-Time Multimedia", Work in Progress.
- [43] Mathis, M., Mahdavi, J., Floyd, S. and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, October 1996.

- [44] Mockapetris, P., "Domain Names -- Concepts and Facilities", STD 13, RFC 1034, November 1987.
- [45] Mockapetris, P., "Domain Names -- Implementation and Specification", STD 13, RFC 1035, November 1987.
- [46] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [47] Partridge, C., Mendez, T. and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.
- [48] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [49] Perkins, C. and P. Calhoun, "AAA Registration Keys for Mobile IP", Work in Progress.
- [50] Perkins, C. and D. Johnson, "Route Optimization in Mobile IP", Work in Progress.
- [51] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.
- [52] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [53] Ramakrishnan, K. and S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP", RFC 2481, January 1999.
- [54] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [55] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [56] Rigney, C., Rubens, A., Simpson, W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2138, April 1997.
- [57] Schulzrinne, H., Casner, S., Fredrick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, January 1996.
- [58] J. Semke, J. Mahdavi, and M. Mathis. Automatic TCP Buffer Tuning. Proceedings of ACM SIGCOMM '98, September 1998.

- [59] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.
- [60] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", Work in Progress.
- [61] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L. and V. Paxson, "Stream Control Transmission Protocol", RFC 2960, October 2000.
- [62] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [63] Touch, J., "TCP Control Block Interdependence", RFC 2140, April 1997.
- [64] Vollbrecht, J., et al., "AAA Authorization Framework", Work in Progress.

A Participants

Juha Ala-Laurila	JUHA.ALA-LAURILA@nokia.com
Mark Allman	mallman@grc.nasa.gov
Alastair Angwin	angwin@uk.ibm.com
N. Asokan	n.asokan@nokia.com
Victor Bahl	bahl@microsoft.com
Fred Baker	fred@cisco.com
Pravin Bhagwat	pravinb@us.ibm.com
Scott Bradner	sob@harvard.edu
Randy Bush	randy@psg.com
Pat Calhoun	Pcalhoun@eng.sun.com
Brian Carpenter	brian@icair.org
Mikael Degermark	micke@cs.arizona.edu
Sally Floyd	floyd@aciri.org
Heikki Hammainen	HEIKKI.HAMMAINEN@NOKIA.COM
Mark Handley	mjh@aciri.org
Bob Hinden	hinden@iprg.nokia.com
Christian Huitema	huitema@microsoft.com
Chih-Lin I	ci@att.com
Van Jacobson	van@packetdesign.com
Phil Karn	Karn@qualcomm.com
John Klensin	Klensin@JCK.com
Jerry Lahti	jerry.lahti@nokia.com
Allison Mankin	mankin@isi.edu
Danny J. Mitzel	mitzel@iprg.nokia.com
Gabriel Montenegro	gab@sun.com
Keith Moore	moore@cs.utk.edu
Eric Nordmark	nordmark@sun.com
Charles E. Perkins	charliep@iprg.nokia.com
Jonne Soininen	jonna.Soininen@nokia.com
Chris A. Wargo	cwargo@cnsw.com
Lars Westberg	Lars.Westberg@era.ericsson.se
Lixia Zhang	lixia@cs.ucla.edu

B Author's Address

Danny J. Mitzel
Nokia
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1 650 625 2037
EMail: mitzel@iprg.nokia.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

