M-Power

A Native Measurement Technique for IPv6-based Networks

Francisco J. Garcia, Robert Gardner & Dimitrios Pezaros

AGILENT LABS

29 January 2004

LANCASTER University Computing Department

Agilent Technologies

Contents

Contents

M-Power

Case

Why IPv6 ?

Case

Address Shortage

- **Success of the Internet being stifled by address shortage**
  - **ALWAYS CONNECTED phones, PDAs, cellular/mobile Systems, sensors of the future will need fixed, constantly available address(es)**
  - **Asia, Europe and Africa currently critically short of addresses**
  - **Class A,B,C system -> uneven distribution of addresses. CIDR not perfect.**
  - **IPv4 Routing tables are exploding resulting in loss of performance.**
  - **NAT extends the life of IPv4 but has serious drawbacks**

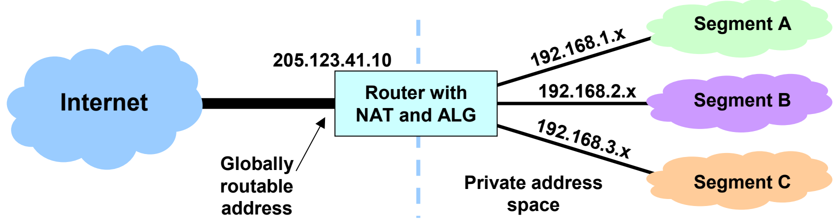

Case

Mobile Communications

- 3GPP Rel. 5: Architecture for All-IP network
  - Radio Access Network: EDGE, UMTS
  - IP Multimedia Subsystem (IMS)
  - GPRS Core Network
  - Mandates use of IPv6 at Application Layer because of address shortage
- 3GPP Rel. 6: Inter-working with WLAN
- Adopting IPv6 throughout:
  - Reduce infrastructure costs
  - Native mobility support (Mobile IP)
  - Native security support (IPSec)

- All-IP Wireless Network characteristics
  - IP-based multimedia services
  - IP-based transport
  - Integration with IETF protocols for functions such as:
    - Wide area mobility (Mobile IP)
    - Signalling (SIP, SCTP)
    - Authentication, authorisation and accounting (Diameter)
  - Network meeting these characteristics is referred to as an *all-IP* network

Case

Wi-Fi and WiMAX

- Carrier class solutions presently poor in areas such as:
  - Manageability
  - Security
  - Mobility
- Carrier class and *low cost* could be facilitated through the adoption of IPv6 features:
  - Connectivity – Stateless node discovery
  - Native security support
  - Native mobility support

Emerging Wireless Standards

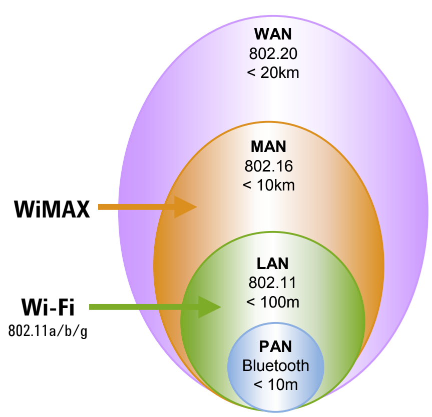

Case

M-Power

Monitoring "Mobile" Services

- When mobile and wireless worlds collide
  - Ubiquitous Internet access with global roaming
  - First commercial deployments of IPv6-based networks
- Devices, users and services are fast becoming mobile
  - Overlay models used to provision and dynamically adapt a delivered service over an existing transport topology
  - Service dissemination topology changes with time
  - Growing trends in Application Level Routing
- Vertical Handovers
  - Access technology changes
    - And so does monitoring infrastructure (e.g. access7 to Wi-Fi sensors)
    - How can we maintain same or similar service assurance functionality?
- Exploit IPv6 enhancements to natively introduce telemetry functionality

M-Power

Approach

A Native Approach

Approach

M-Power

Accessional Techniques

- Issues with Passive Measurements
  - Scalability –probes, correlation engines, data volume etc.
  - Link monitoring has its limits
    - New service or new link could trigger the need for complex re-engineering of link monitoring hardware
    - Growing demands for high capacity links of 10Gb/s or higher
  - Complexity involved in performing 2-point measurements
  - Challenging to infer end-to-end view
  - Further complicated with introduction of security
- Issues with Active Measurements
  - Assume measured performance for synthetic traffic reflects performance of user traffic
  - Measurement traffic itself may be a factor in performance degradation
  - Measurement mechanism tightly coupled with the measurement applications

Approach

M-Power

Extension Headers

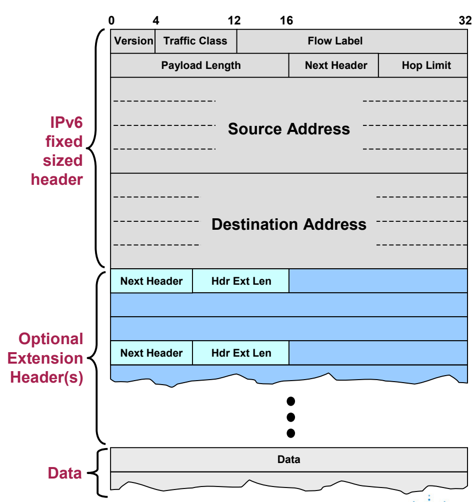

- IPv6 Extension Headers
  - Optional information encoded in separate headers between IPv6 header and upper-layer header.
  - Packet may carry zero or more EHs.
  - So far there are only a handful of standardised EHs:
    - Hop-by-Hop Options.
    - Routing (Type 0).
    - Fragment.
    - Destination Options.
    - Authentication.
    - Encapsulating Security Payload.
  - With one exception, EHs are only processed at the destination(s).
  - The exception is the Hop-by-Hop Options, processed at every node.

Approach

Extension Headers and Telemetry

- **Exploit extension headers to introduce *native* measurement and management functionality**
- **Applying these notions to the instrumentation of measurements**
  - **In-line measurements –piggybacking triggers and measurement data onto real user traffic**
  - **Lowest level condition-event-action triggers for influencing measurements system behaviour**
  - **Multipoint measurement technique**
- **Examples using destination header options**
  - **One-way loss**
  - **One-way delay**

TLV-encoded Options

Option Type	Option Length	Data
1 octet	1 octet	variable number of octets

Destination Options Header

Next Header	Hdr Ext Len	
	Options	

Timestamp Destination Options TLV

		Option type	Option data len
Pointer	Overflow	Flags	(Reserved)
Source timestamp: seconds			
Source timestamp: microseconds			
Destination timestamp: seconds			
Destination timestamp: microseconds			

Loss Destination Options TLV

	Option type	Option data
Sequence Number		

Approach

Applicability

- **A native approach for introducing service measurements**
  - **Not intended as a replacement for active/passive techniques**
  - **Complementary technique, when to use depends on service characteristics being monitored (e.g. *mobile* services)**
- **Facilitates seamless and incremental deployment**
  - **Distribute telemetry intelligence to location where it is required, when it is required**
  - **Allows for the engineering of distributed monitoring solutions that dynamically adapt to the monitored service**

Approach

Comparative Analysis

M-Power

Aspect/Property	Active Measurements	Passive Measurements	Inline Measurements
Impact on network (Measurement process)	- Intrusive: Generates additional load which competes for resources	++ Non-intrusive: No impact on network	+ Intrusive: Marginal load increase and minor delay might be incurred
Impact on network (Measurement data)	+ Load generated at one end point	- Load generated at one or both ends	+ Load generated at one end point
Confidence	- Artificially injected traffic used to infer/predict experience of real traffic - Test traffic may be treated differently - Injected traffic affects performance	+ Measures real user traffic	+ Measures real user traffic - Possibility that instrumented traffic is distinguishable and treated differently
Controllability	+ Can test any traffic, path, method of sampling, protocol, etc. – at any time.	- Can only measure available traffic	- Can only measure available traffic - Requires an accommodating protocol
Security/Privacy issues	+ Private, injected traffic + Real data not examined	- Observing real traffic	-- Observation and modification of real traffic
Scalability issues	+ Can be dynamically deployed on a per interface basis + Can inject a chosen amount of traffic	- Probes per interface at ingress & egress - Full packet capture is not scalable + Can use filtering and sampling	+ Can be dynamically deployed on a per node or per interface basis + Can use filtering and sampling
Complexity and Processing	+ Correlation not required - Non-trivial generation of statistically representative test patterns	- Correlation of large quantities of data from ingress and egress is computationally intensive and doesn't scale well	+ No correlation - Statistical sampling and filtering
Major application areas	Two-point measurements: Quality of Service testing, such as available bandwidth, trip delay, and packet loss.	One-point measurements: packet filtering and counting to obtain traffic type, source / destination etc.	Multi-point, policy-based measurements, active troubleshooting, packet loss, delay, tracing, routing, packet / flow foot printing.
Other comments	- Eavesdropping not possible - Requires substantial expertise to produce meaningful test patterns	+ Eavesdropping possible	+ Eavesdropping possible - Not applicable to all traffic types (e.g. real-time, max MTU traffic)

Approach

Ubiquitous Measurements

M-Power

- The need for ubiquitous measurements is ever growing and this is exemplified by current IETF activities
  - Passive Sampling – standard set of capabilities for network elements to select subsets of packets by statistical and other methods that may assist in baseline measurements, performance measurements, troubleshooting., etc. http://www.ietf.org/html.charters/psamp-charter.html
  - The case for an Internet Measurement Protocol – allowing it to be handled by the forwarding path rather than the router CPU. http://www.irtf.org/charters/imrg.html
  - IP Flow Information Export – network elements exporting flow information in a standard way so that it can be fed directly into mediation, accounting/billing and network management systems. http://www.ietf.org/html.charters/ipfix-charter.html
  - IP Performance Metrics – defining standard set of metrics and active measurement procedures to accurately measure such metrics. http://www.ietf.org/html.charters/ippm-charter.html
- Are these approaches flexible, adaptable and scalable enough to handle *mobility*?
- Issues
  - Standard processes could be lengthy
  - Potential for lots of non-service specific data
  - New service, new set of metrics – potentially long time before the measurements could be deployed

Prototype

Functional Prototype

RIPE 47 Meeting, Amsterdam
29 January 2004

Agilent Technologies

Slide 15

Prototype

Applying Programmable Networking Concepts

- Dynamically programmable network architectures
  - Supporting fast service creation and deployment
  - "network aware applications and application aware networks"
- Applying these concepts to the telemetry of *mobile* services
- Other lessons could be learnt from peer-to-peer, ubiquitous and pervasive computing
  - Deploy agents to perform specific functionality
  - Exploit user equipment
  - Agent-based approaches facilitate immediate deployment
- Telemetry Agents
  - Lowest level deployable component
  - Remotely managed
  - Dynamically linked into application
- Advantages
  - Transparently introduced
  - Dynamically deployed
  - Engenders flexibility, extendibility and scalability

RIPE 47 Meeting, Amsterdam
29 January 2004

Agilent Technologies

Slide 16

Prototype

Test-Beds

- Test-Beds at Lancaster and Agilent Labs are based predominantly on Linux systems (Kernel 2.4.x)
  - Telemetry agents implemented as Loadable Kernel Modules (LKM)
    - Can be linked with a running kernel at run-time
  - Distributed control, security, communications and management frameworks built using Java technologies
    - Including streaming of service data records
  - Test-beds consist of a combination of PCs to act as servers, wireless access points (802.11b) and routers
  - Mobile nodes based on laptops, Sharp Zaurus, and iPAQ
  - Mobile IPv6 used for managing mobility

Prototype

General Architecture

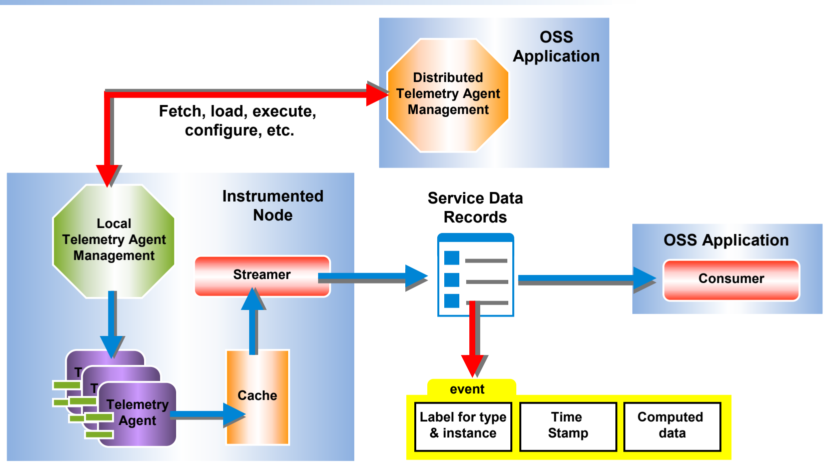

Prototype

Adding Telemetry Headers

M-Power

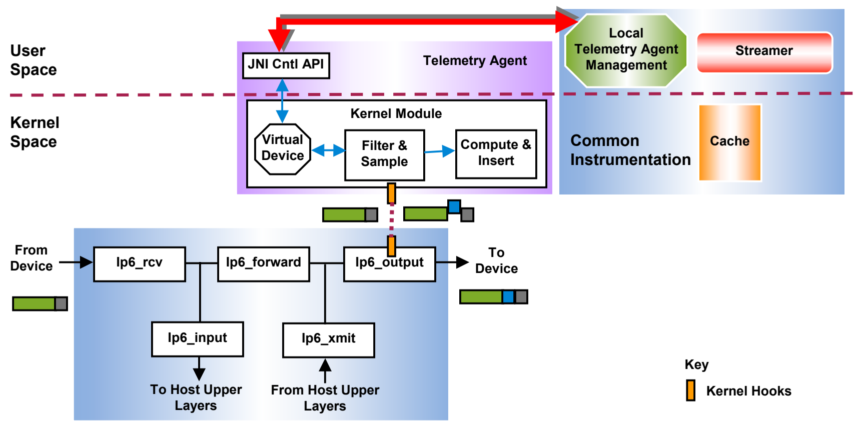

Prototype

Removing Telemetry Headers

M-Power

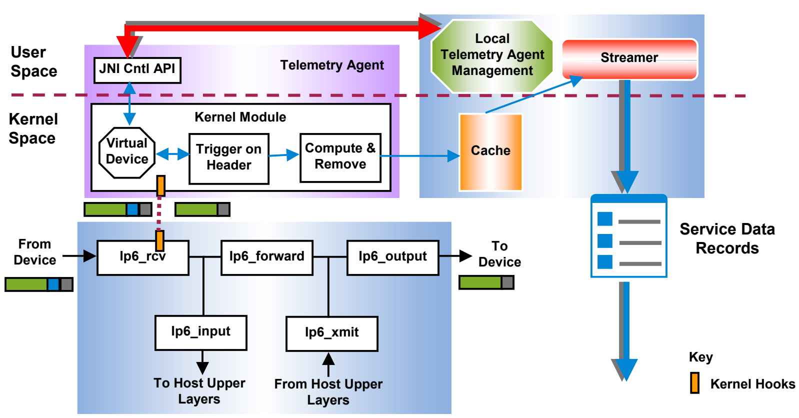

Results

Experimental Results

RIPE 47 Meeting, Amsterdam
29 January 2004

Agilent Technologies

Slide 21

Approach

Example Measurements

- These are very basic measurements over our own "small" test-beds
  - Not very exciting, involves running our own applications and injecting our own traffic
  - Applications include:
    - Video streaming
    - Interactive TCP (Telnet, SSH)
    - Web browsing
    - Bulk TCP transfers
- Initial experiments applied end-to-end
  - On all traffic
  - Filtered and sampled traffic
  - To evaluate the efficacy of this native measurement approach
- Recently started looking at applying the technique to signaling protocols
  - SIP
  - Mobile IPv6

RIPE 47 Meeting, Amsterdam
29 January 2004

Agilent Technologies

Slide 22

Approach

TCP Goodput

M-Power

Service Port	Client Port	Conv. Setup Time (msec)	Conv. Duration (sec)	Packets	Completeness
21	32809	49,605	504	111	True
51694	32810	52,665	0.0	8	True
8320	32811	48,221	78.0	4831	True
37760	32812	60,588	-1.0	4295	True
41782	32813	62,778	0.0	8	True
14253	32814	50,731	139	4943	True
53059	32815	57,880	121	3103	True

*Blue: FTP Control Channels
*Black: Data transfer from server to client (GET)
*Red: Data transfers from client to server (PUT)

Graph for MGET
X-points: Data packets from the server
Red Dots: Acks from the client

Example results obtained over operational broadband network through ADSL Connection

Graph for MPUT
X-points: Acks from the server
Red Dots: Data packets from the client

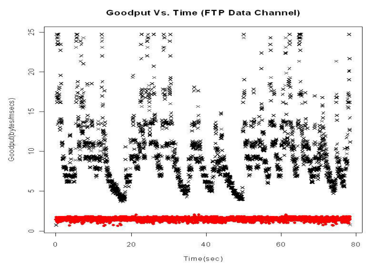

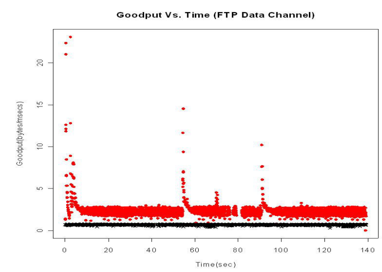

Approach

Video Streaming

M-Power

Min: 8 / Max: 89 / Ave: 20.9 (msec)

Min: -68 / Max: 58 / Ave: -0.00049

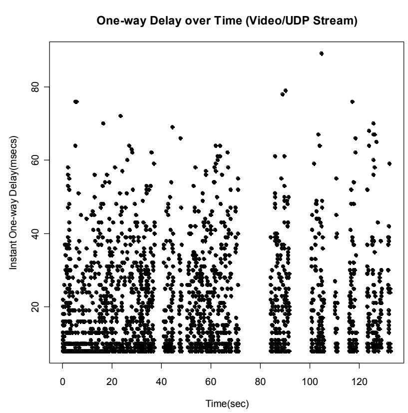

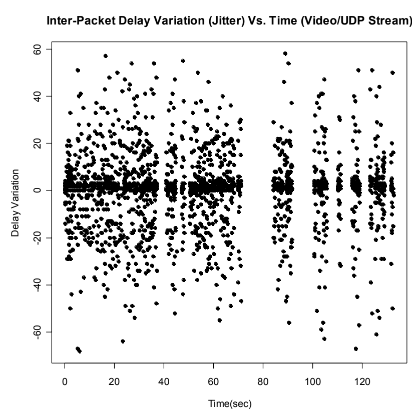

Approach

SIP Signalling

M-Power

- **Applied to**
  - SIP call set-up
  - SIP client registration
- **Example telemetry agents developed**
  - Sip_Delay: application-agnostic agents adding simple timestamps to all SIP/UDP messages sent/received
  - Sip_Filter: stateless application-aware agent that can be configured at runtime to filter on specific SIP/UDP messages (e.g. INVITE)
  - Sip_Register_Time: stateful application-aware agents applied to the measurement of SIP client registration with register proxies.

Example SIP Registration Results

Client	Registrar	Registrar Time	Transit Time	Total Time
<sip:evbouzianis@ismene-6>;tag=3173795502	ismene-6	1205	594	1799
<sip:evbouzianis@ismene-6>;tag=1448846863	ismene-6	2237	562	2799
<sip:evbouzianis@ismene-6>;tag=4225653434	ismene-6	2652	580	3232
<sip:evbouzianis@ismene-6>;tag=1751224854	ismene-6	1054	631	1685
<sip:antigone@ismene-6>;tag=182395451	ismene-6	1118	1065	2183
<sip:antigone@ismene-6>;tag=3514547268	ismene-6	1145	1383	2528
<sip:antigone@ismene-6>;tag=998640064	ismene-6	1030	1837	2867

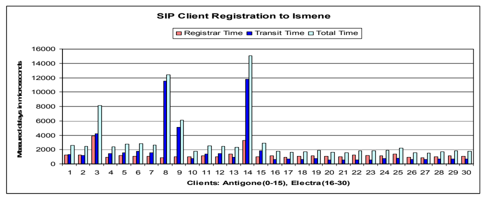

M-Power

Future

Next Steps and Concluding Remarks

Future

Next Steps & Concluding Remarks

M-Power

- Study and quantify the performance/cost in deployment and operation of this scheme
- Engage with interested parties
  - Larger trials on operational or test-bed IPv6 networks
- Evolve SIP and mobile service monitoring
- Study other application domains

Acknowledgements

M-Power

- Dimitrios Pezaros from Lancaster University, is an Agilent sponsored PhD student studying the efficacy and applicability of the presented native approach to telemetry
- Evangelos Bouzianis from Lancaster University, was sponsored by Agilent during his MSc dissertation and developed the SIP monitoring capabilities (http://ww.lancs.ac.uk/postgrad/bouziani/)
- Prof David Hutchison from Lancaster University, is Dimitrios academic supervisor and has supported him on this work and provided continual and valuable input.
- Prof Joe Sventek, was formerly head of Agilent Research Labs in Scotland before accepting his current chair with Glasgow University. This project was initiated under his supervision and he has again provided continual support and valuable input.