

Quality of Experience

QoE is not just about speed, but more about the other factors that impact our ability to deliver great video, browsing and gaming experiences

Philip Eardley,
Trevor Burbridge, Arnaud Jacquet, Alan Smith, Andrea Soppera, Achilles Petras

30th October 2015

RAIM workshop

BT

Leone

From global measurements to local management

Philip Eardley (BT) – Coordinator

Nov 2012 – April 2015

www.leone-project.eu

The research leading to these results has received funding from the European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 317647.

Saba Ahsan, Jorg Ott (Aalto) - YouTube
Magnus Boye, Alemnew Asrese, Pasi Sarolahti (Aalto) – web rendering
Boris Banjanin (MG-Soft), Prapa Rattadilok (RGU) – auto data analysis
Sam Crawford (all tests)
Marcelo Bagnulo (UC3M), Juergen Schoenwaelder, Vaibhav Bajpai (Jacobs), Trevor Burbridge (BT) - standards

Measuring Quality of Experience

Monitor and study broadband demand behaviour and performance

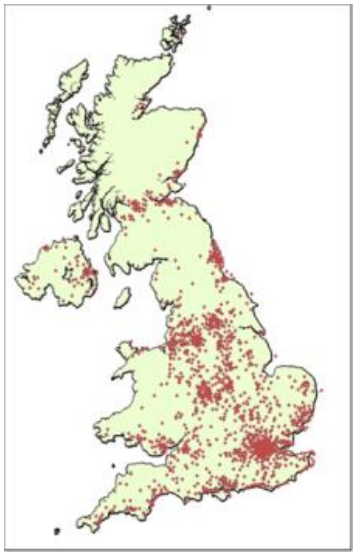

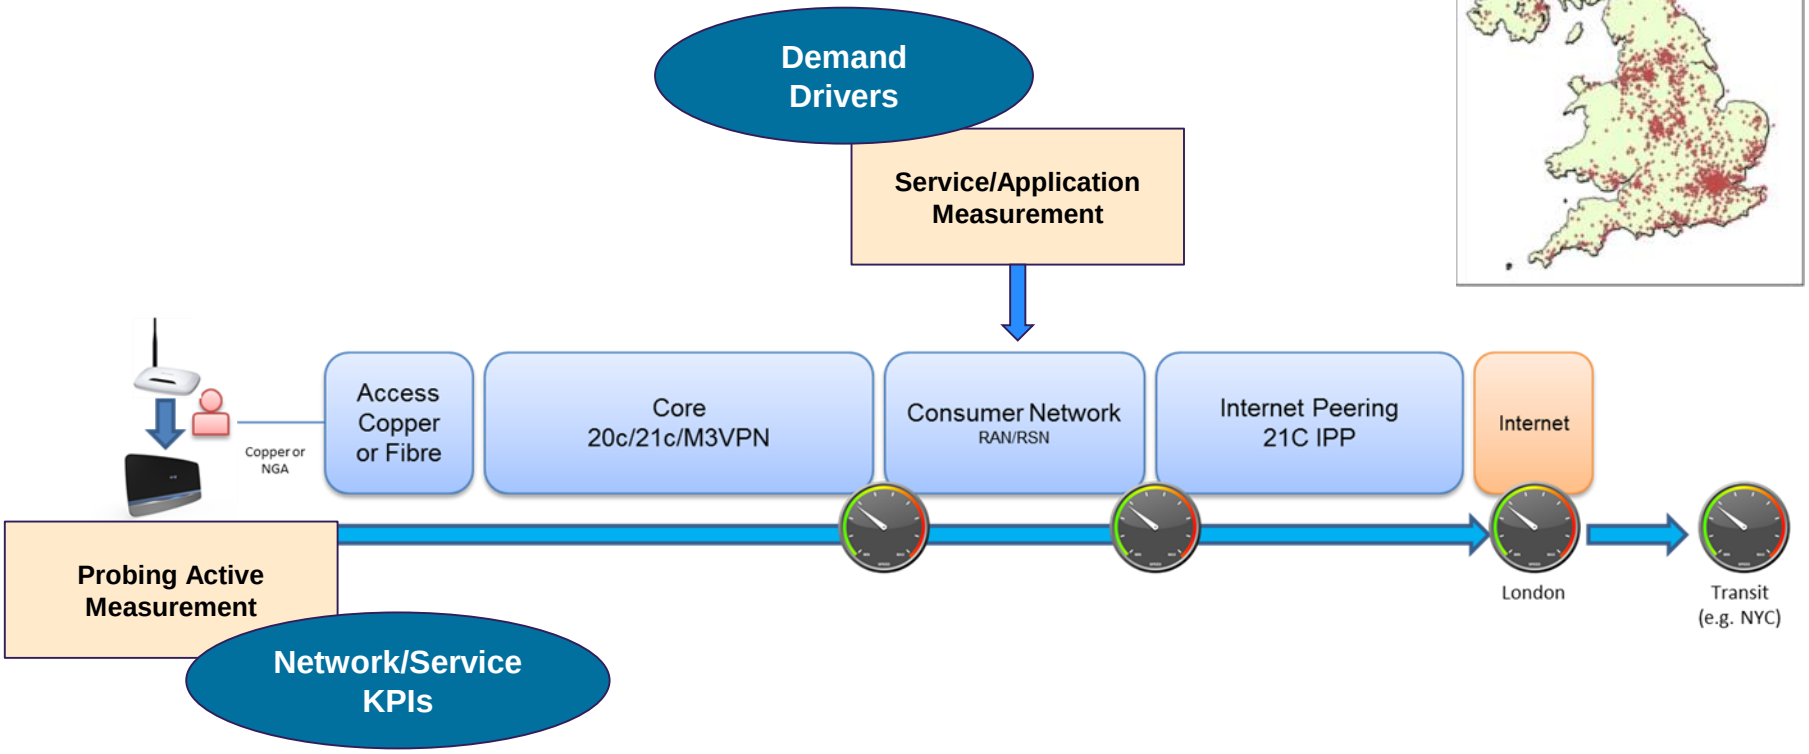

- **Passive measurements –** per-line usage statistics
- **Active measurements -** set of tests (speed, packet loss….) run on selected lines.

Historic traffic growth observed on Broadband

BB Peak Time Gbit/s view (last 10 years)

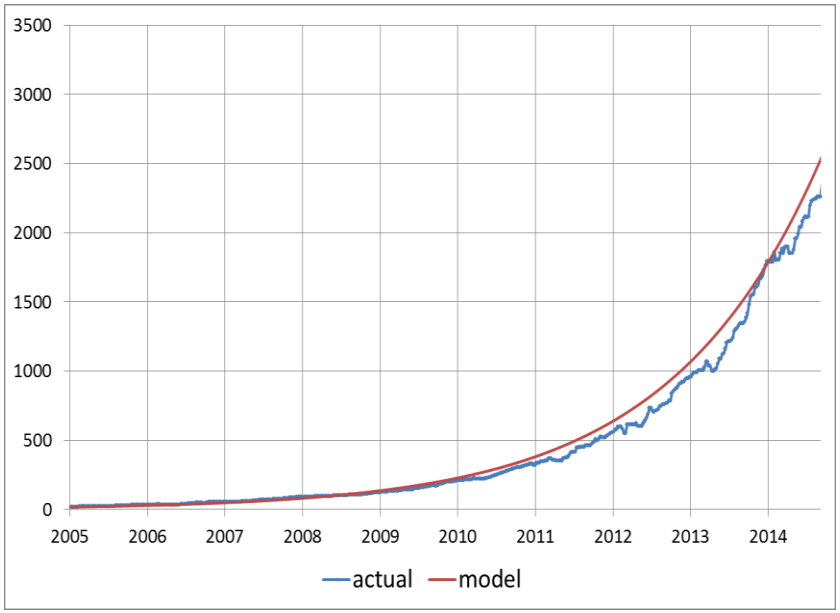

BB Peak Time Gbit/s view (Log Axis)

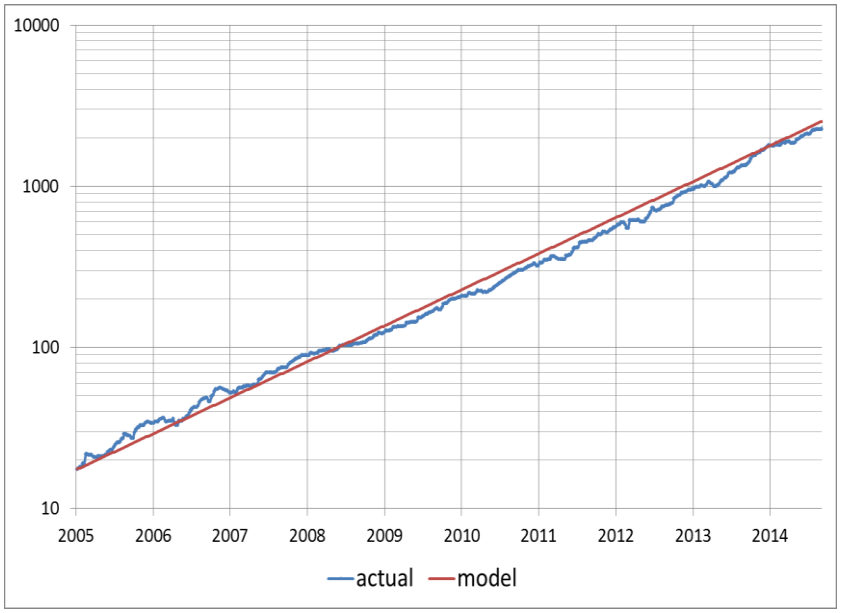

- Total network demand has grown more than 100 times over last ten years
- Core broadband traffic grows at 65%+ year on year growth
- Driven by: video (already 60% of total demand) and evolution of access
- Note – this is just broadband traffic – excludes all business and other services

To be published: The Impact of Capacity Growth in National Telecommunications Networks
AndrewLord*, Andrea Soppera, Arnaud Jacquet. Phil. Trans. R. Soc. A.

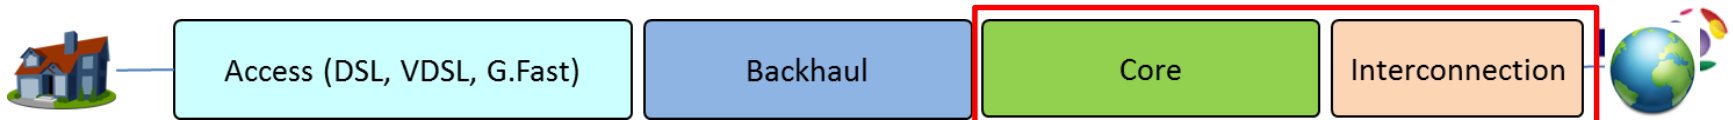

Large-scale active measurements – helping us to handle network growth

- Identify hotspots in the network
  - At some level of aggregation
  - Understand impact on user's experience
- Understanding the impact and operation of new devices, technology, products and services
  - Caching to mitigate growth
  - IPv6, IPTV, Home Gateways, new line cards…
- Other ISP use cases
  - Identifying and isolating failures in network
  - Identifying issues on an individual line
  - To monitor suppliers (upstream & downstream)
  - Understanding customer's end-to-end service experience (e.g. web browsing quality; reliability)
- Also regulator and end-user use cases

Measuring Quality of Experience

- Active reference testing
  - able to accurately correlate & detect problems
- End to end
  - pick up any problems at any point/layer
- User experience
  - assess service & user impact

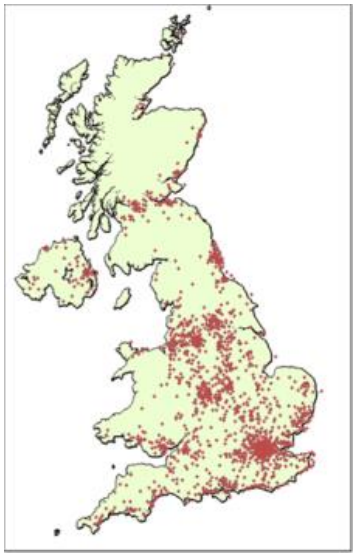

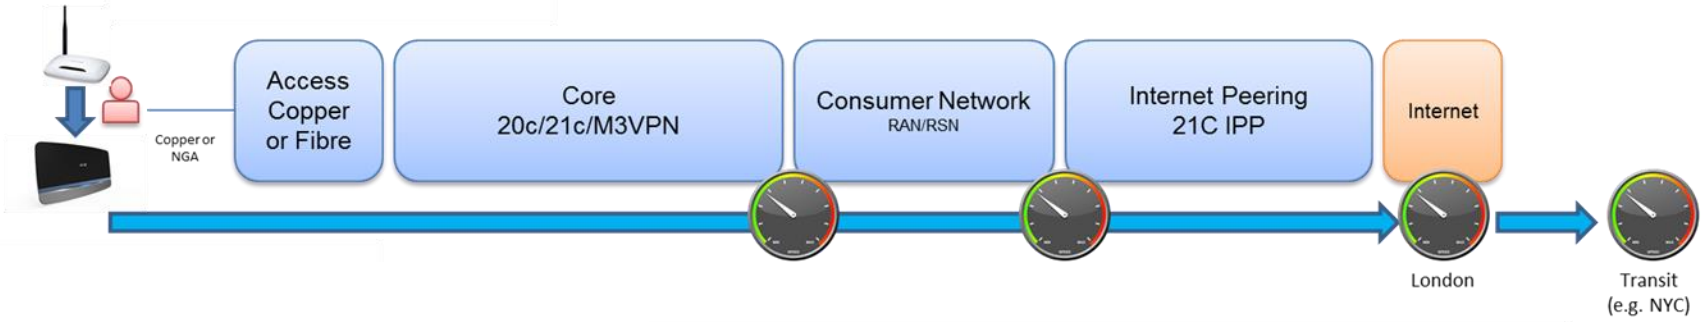

Portal Overview

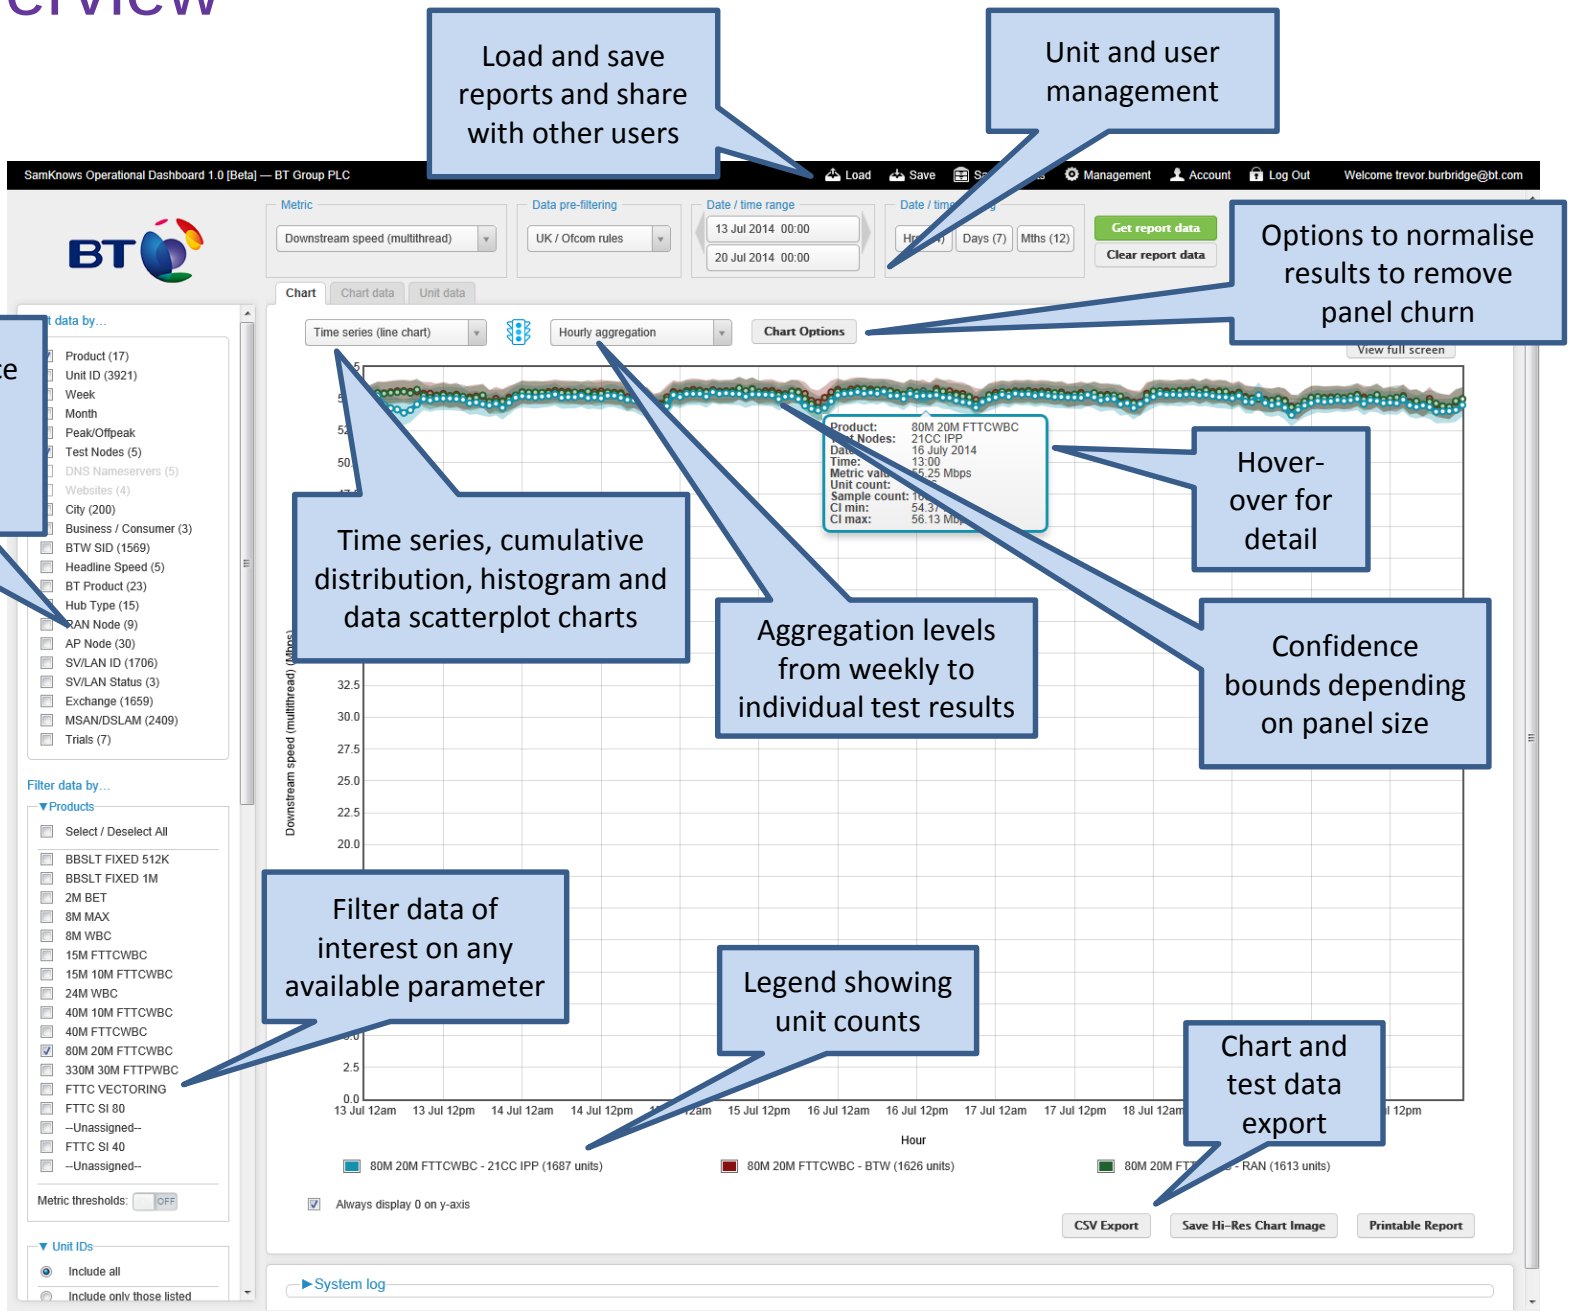

Commonly-used Charts

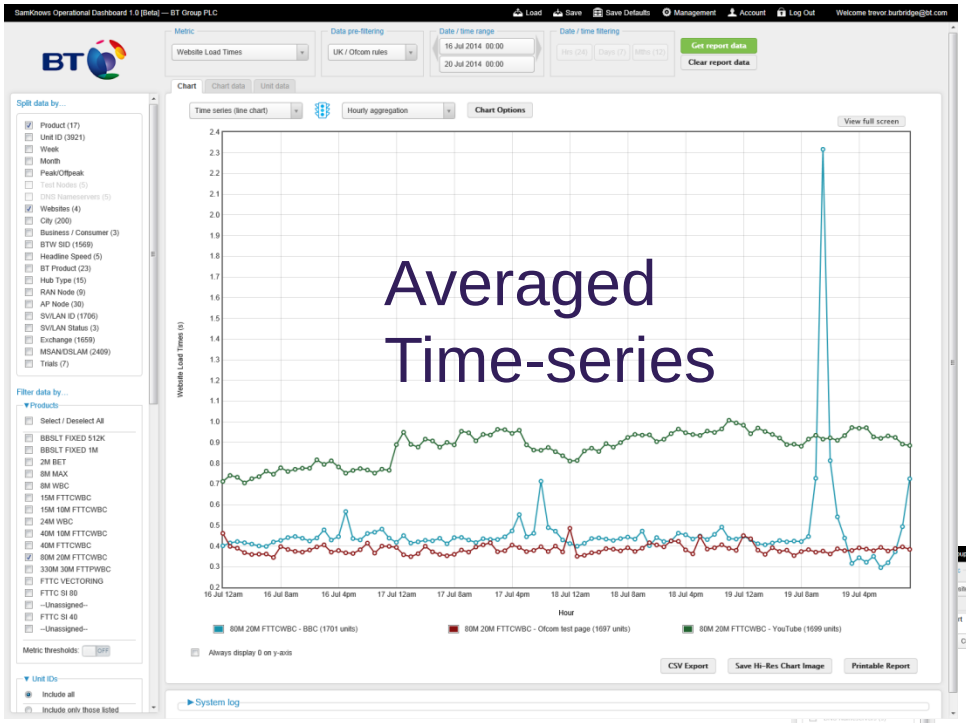

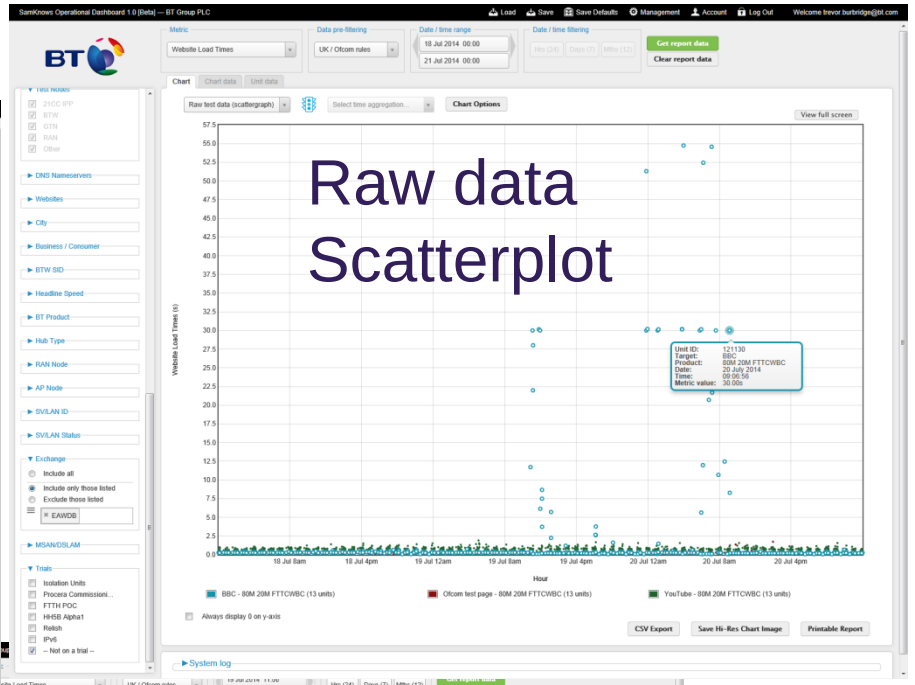

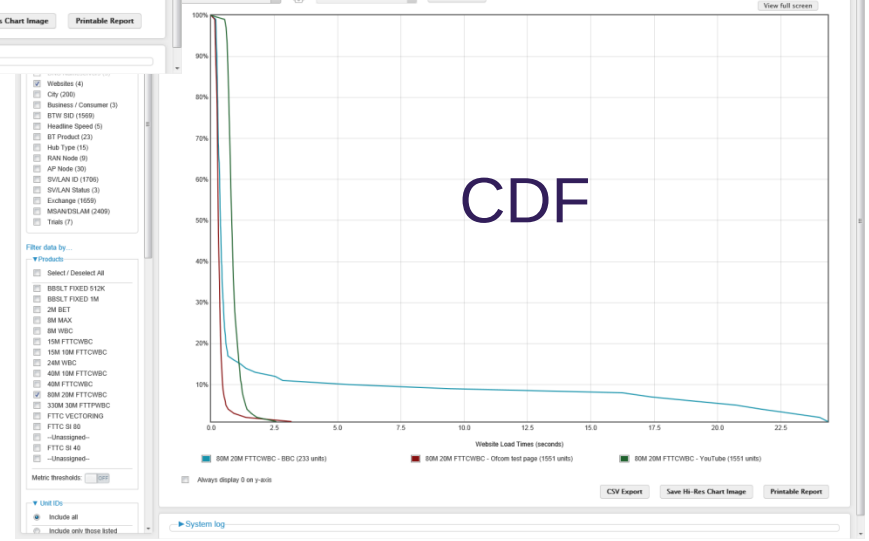

iPlayer and caching

- Catch-up for BBC programmes
- How does caching work and how well?

iPlayer and caching

- iPlayer content comes at several characteristic rates, the most dominant being 2.8Mbps, 1.5Mbps and 0.8Mps
- three CDNs are used
  - “a” CDN only hosts 2.8Mbps
  - “c” CDN doesn’t host 2.8Mbps
- XML manifest assigns a priority
  - ‘fast’ lines “a” or “b” 50:50 basis
  - ‘slow’ lines “b” or “b” 50:50 basis
- (Top pic) “a” and “b” have different start-up delays due to different source rate limit
- (lower pic) Test reported drops in reliably streamed bit rate (in red), due to failures on “a” CDN (in blue)
- Note: iPlayer & caching has changed recently

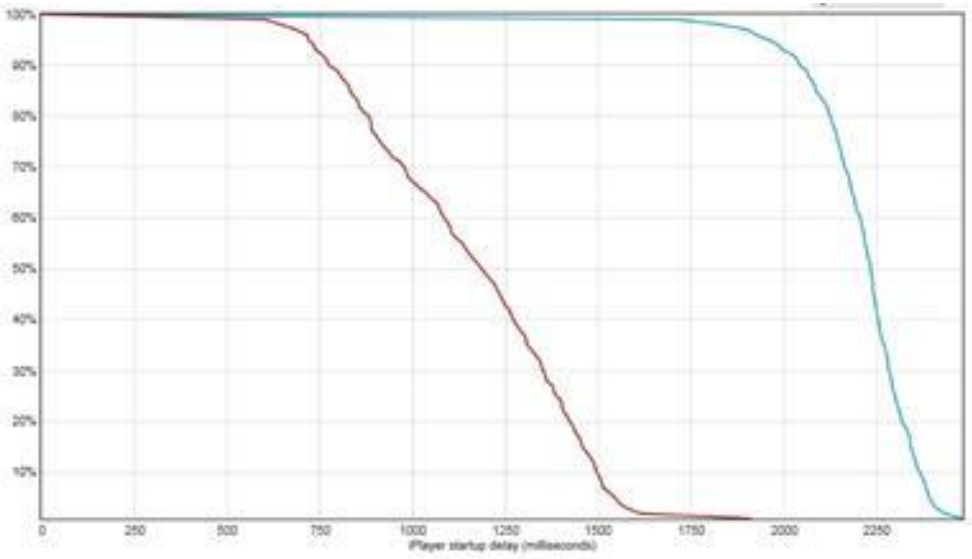

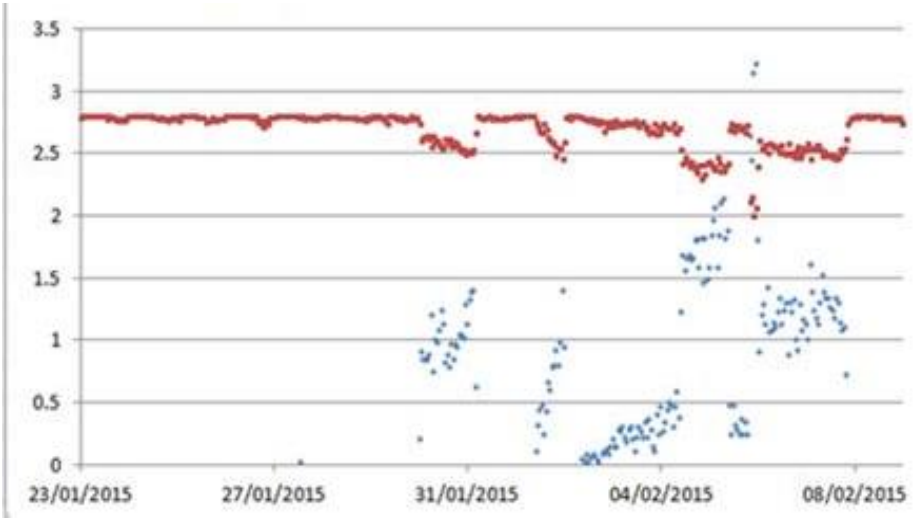

Web rendering test

- TCP download time may not accurately reflect user experience
  - QoE OK when first 80% of visible content downloaded?
- Test looks every 100ms to see if pixels changed on the browser screen – complete is no change for 3secs

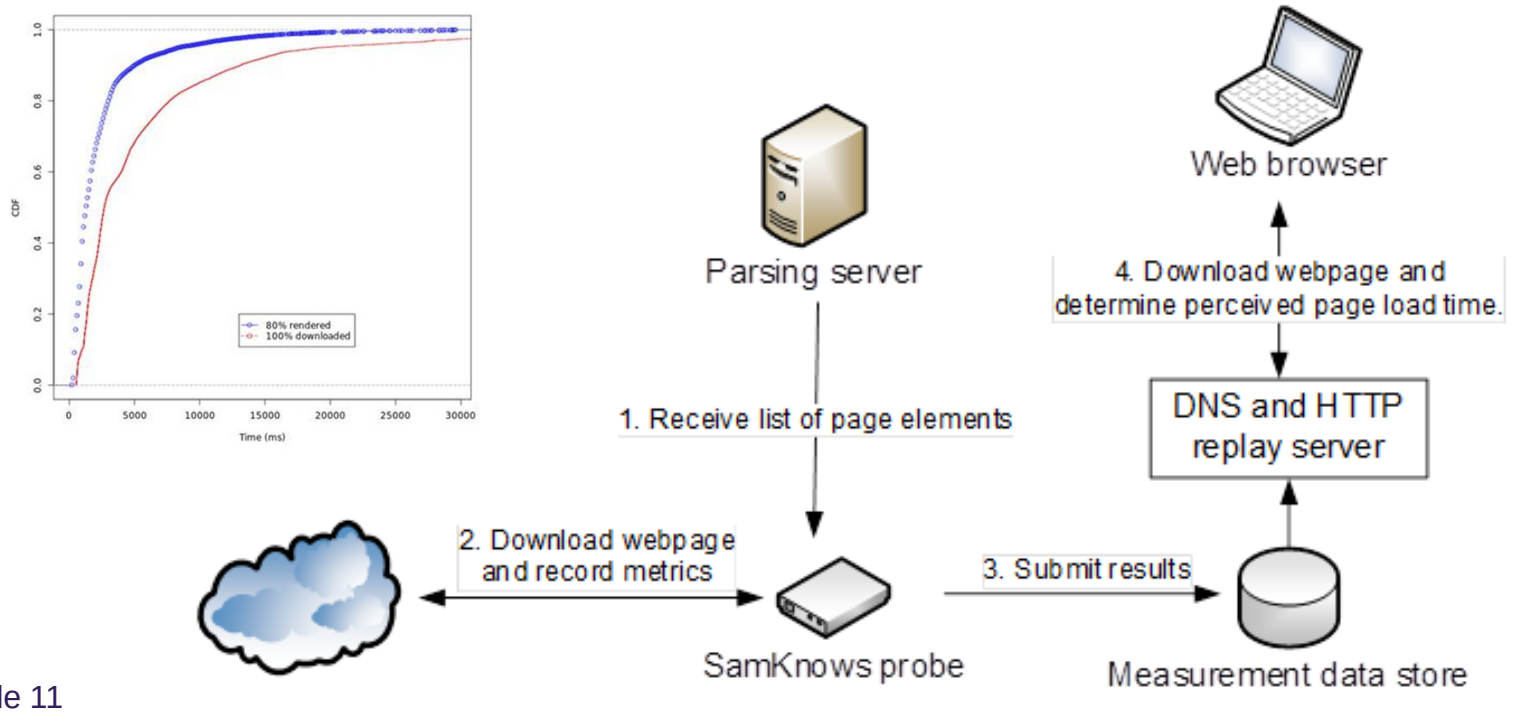

Web rendering test - results

- Correlation of rendering time with ping (left) & throughput (right)

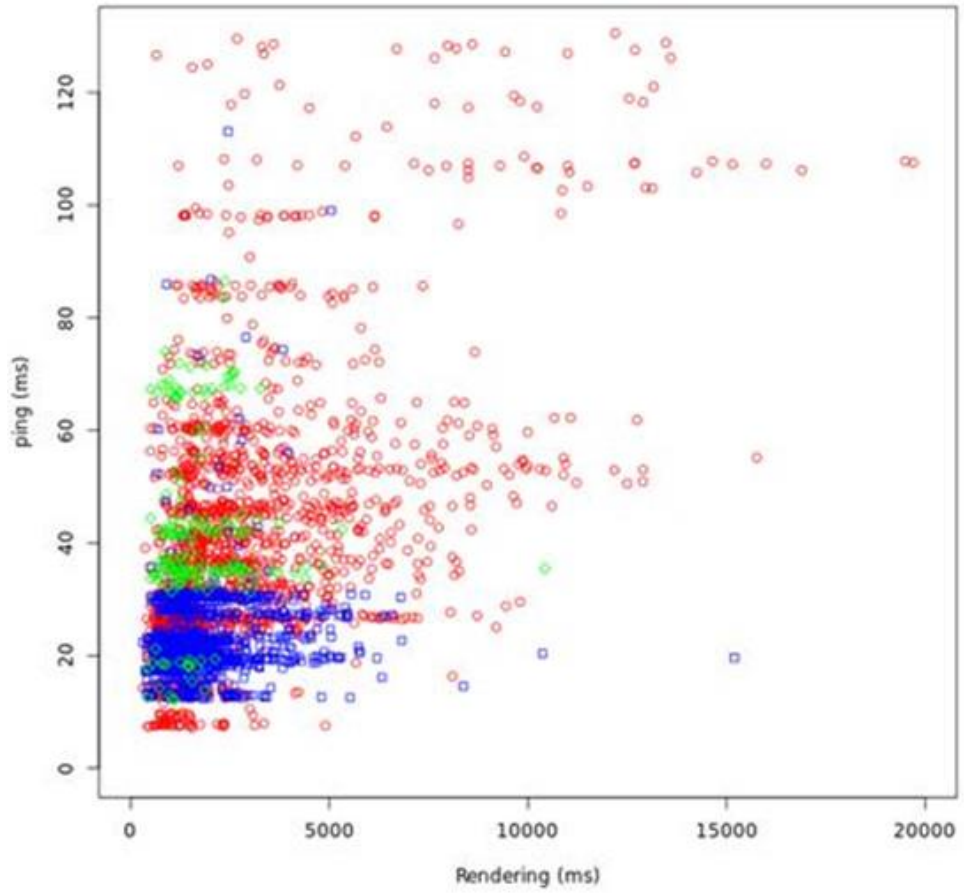

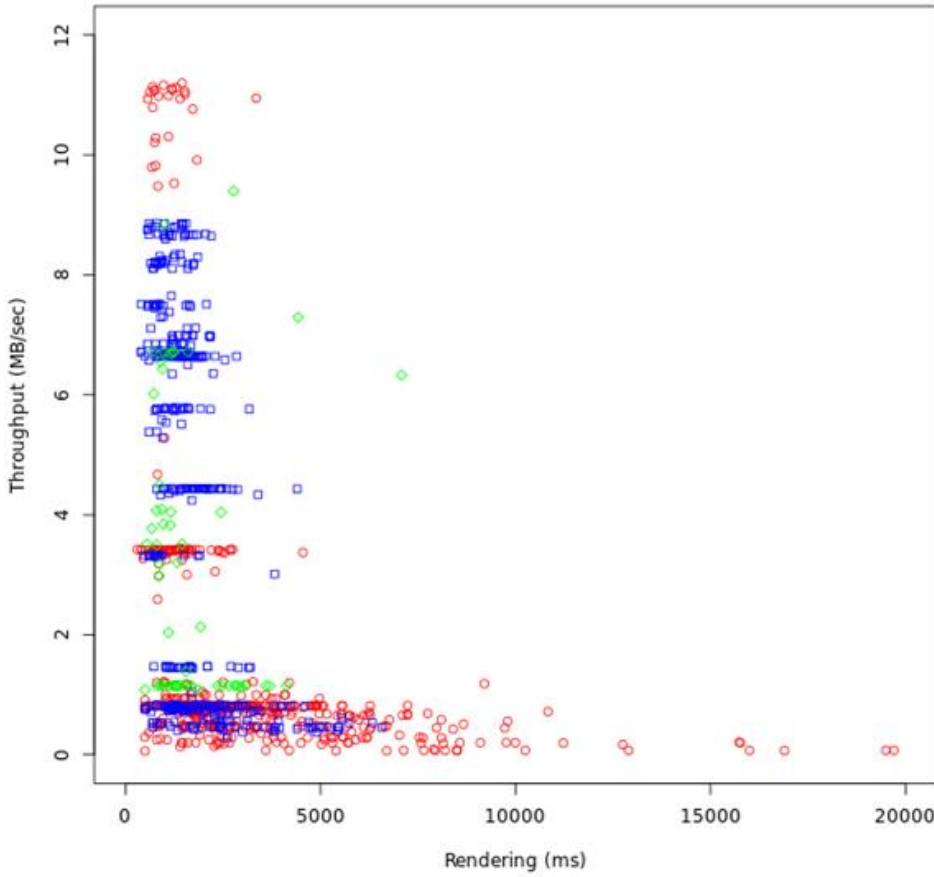

Some opinions

- More realistic tests (video, VoIP)
- Schedule – hourly about right
- Metadata inaccuracies – tests to check
- Data cleansing – eg outages impact pkt loss
- On-net servers
- Benefit from identifying shared issues
- Per-line potential benefit

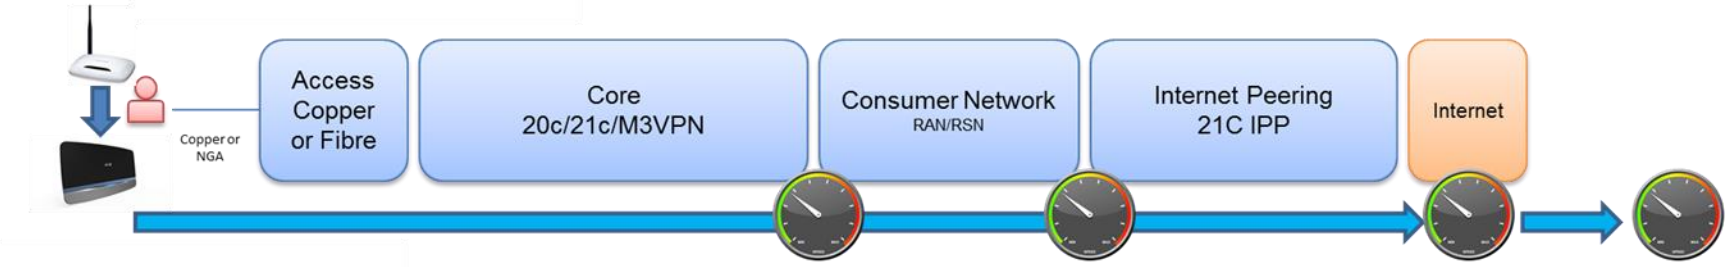

Missing pieces & Research areas

- Finer granularity needs more probes
  - From hardware to software
  - Big stop button
- (Automated) Data analysis
  - New tools to scale performance and improve usability (big data)
  - On-demand testing (call centre)
- Improved Diagnostics
  - Available capacity testing
  - Identifying problems in the home network
  - Supply chain analysis
- Standardisation
  - Meaningful to compare measurements of same metric
  - Allow operators to use multiple vendors

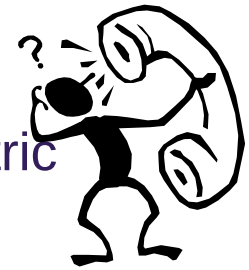

Automated data analysis

- Motivation: identify sudden failures, long-term degradation…
- Assistance to network manager: Goldilocks number of alarms
- Open questions
  - Real-time?
  - Training history in /out?
  - Multiple metrics?
  - Accuracy?

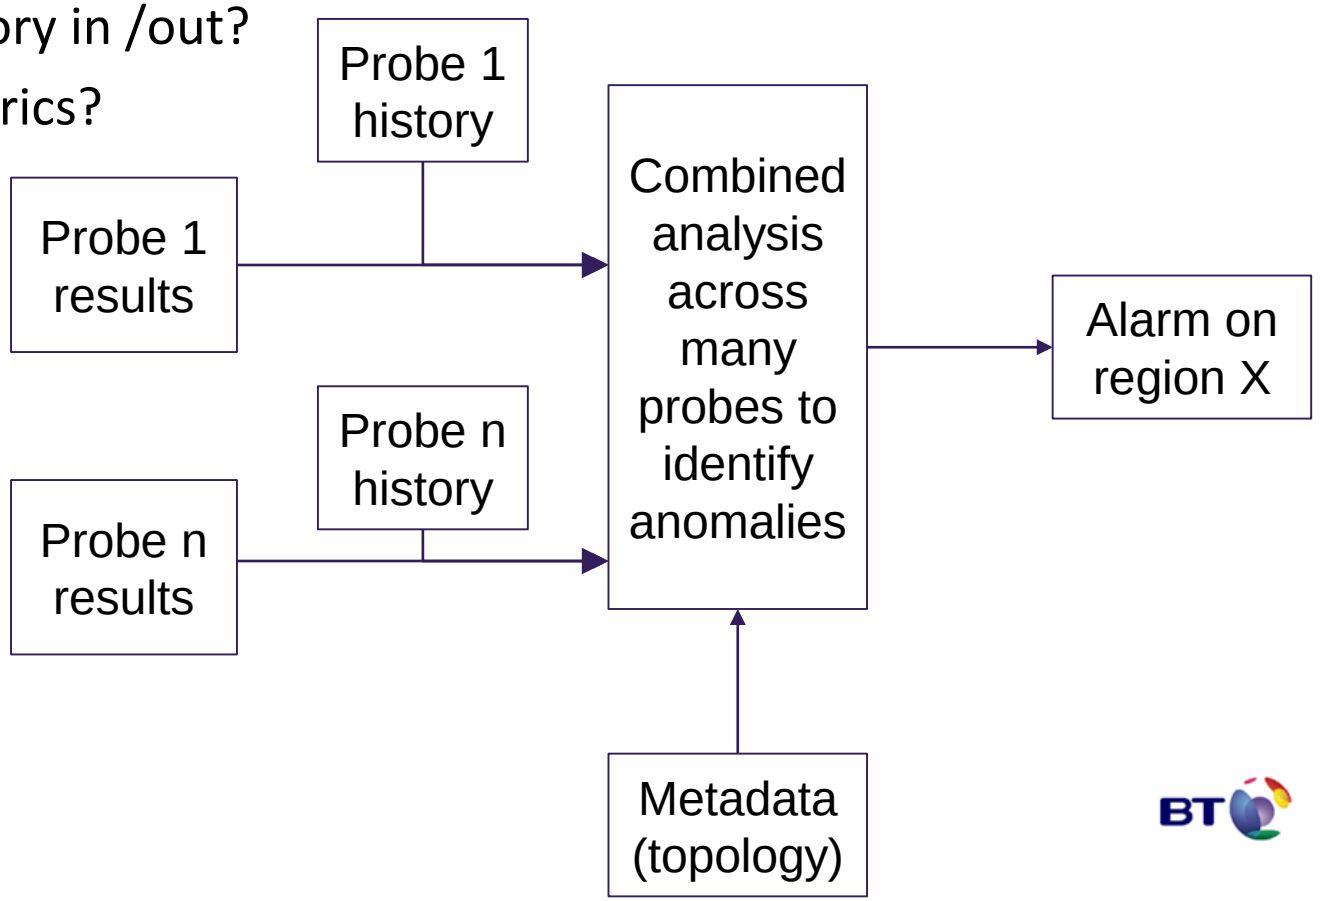

Capacity Testing

- Running Throughput tests on many lines is heavy on the network and potentially ties up user lines (even for a few seconds)
  - Too few probes cannot give good visibility of capacity problems in the SVLAN/VP
- Solution: use large number of hubs with lightweight capacity tests
- Basic principle: send short packet trains (or pairs) into the network and analyse dispersion

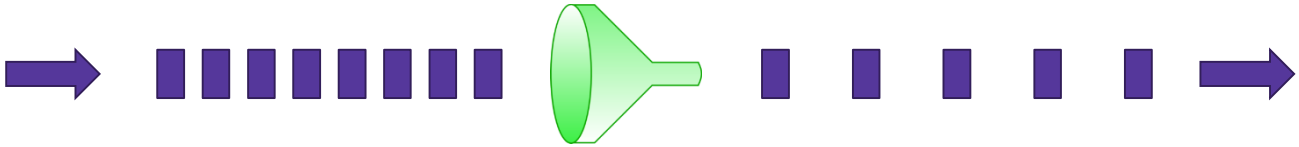

- Different tests to detect capacity vs. available bandwidth
- Approaches
  - Packet pairs vs trains
  - Iterative vs. direct probing
- Overcome accuracy problems from multi-hop delays
- Don't want to affect other traffic
- But do want to see impact of other traffic

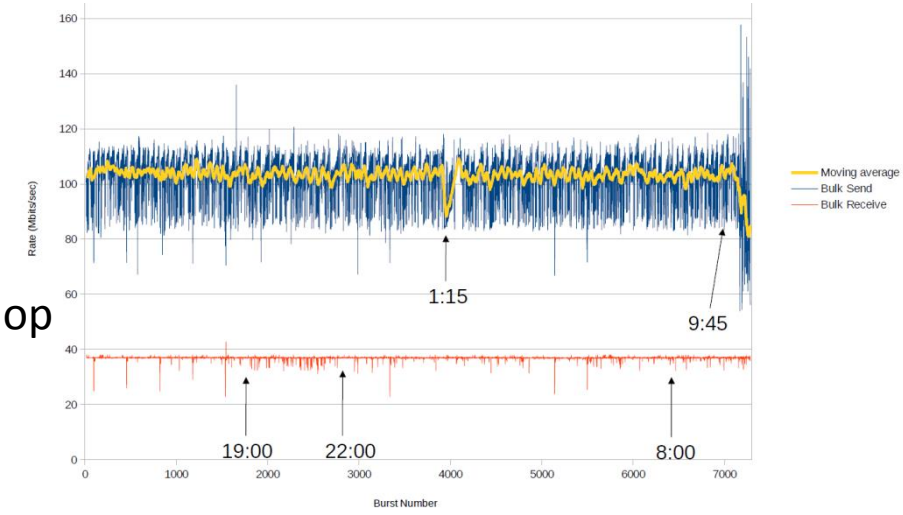

Home Network Testing

- Self-help tool for customers
- ISP wants additional insight into home network and device performance
- Use lightweight probe-based techniques such as traceroute and device discovery?
- Passive analysis of devices connecting through home gateway?
- Install on user device?
  - Single viewpoint limited
  - Forced user participation

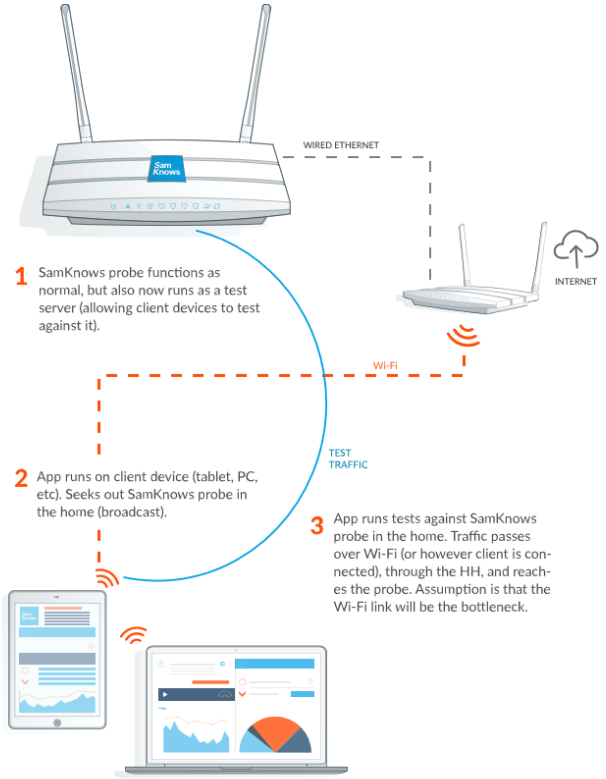

Supply Chain Mapping

- Try to detect where problems are in the network between users and the global services they access
- Not limited to BT on-net but gain a view of global routes, especially to popular services, and also home network
- Helps diagnose service problems and negotiate better peering and transit arrangements

Supply Chain Mapping – use Traceroute?

- Possible approach: probe delay to each ‘hop’ along the path to a range of destinations
  - Look at daily increase in delay variation
- Looking at overall delay variation can fall foul of equipment that has variable response to replying to traceroute TTL expiry
  - Ie ‘problems’ may not affect normal traffic
- How to filter out misleading data?
- High delays and variation in early hops can mean later hop delays can be hidden in the noise
  - Since each hop probe is separate packet
  - Essential to have quiet line or what you will measure is simply impact of user traffic on their own line
- Would be nice to have ping++ !

CAIDA: Archipelago (MIT)

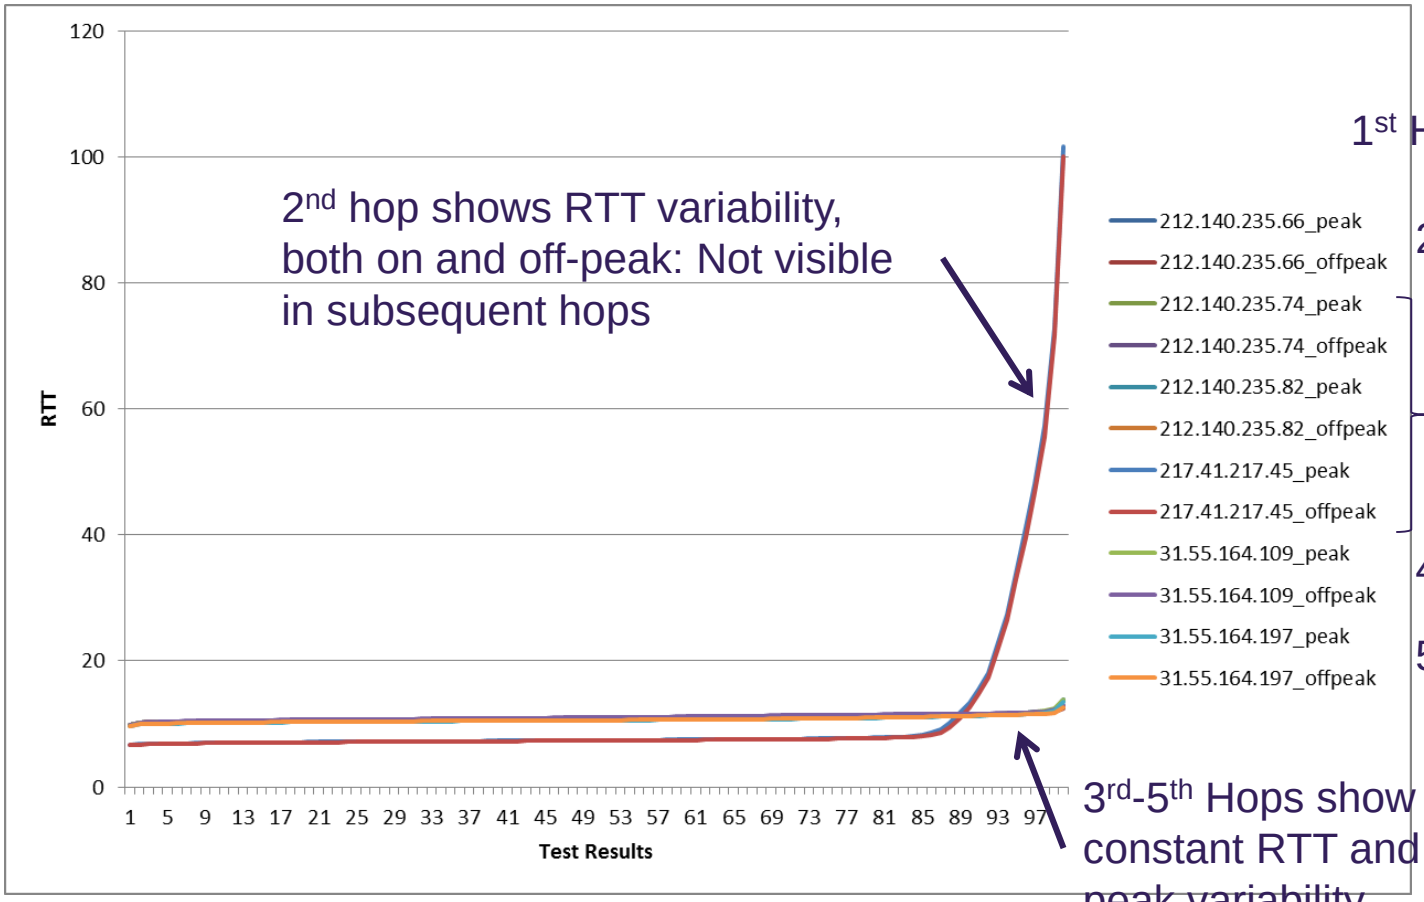

Standards perspective

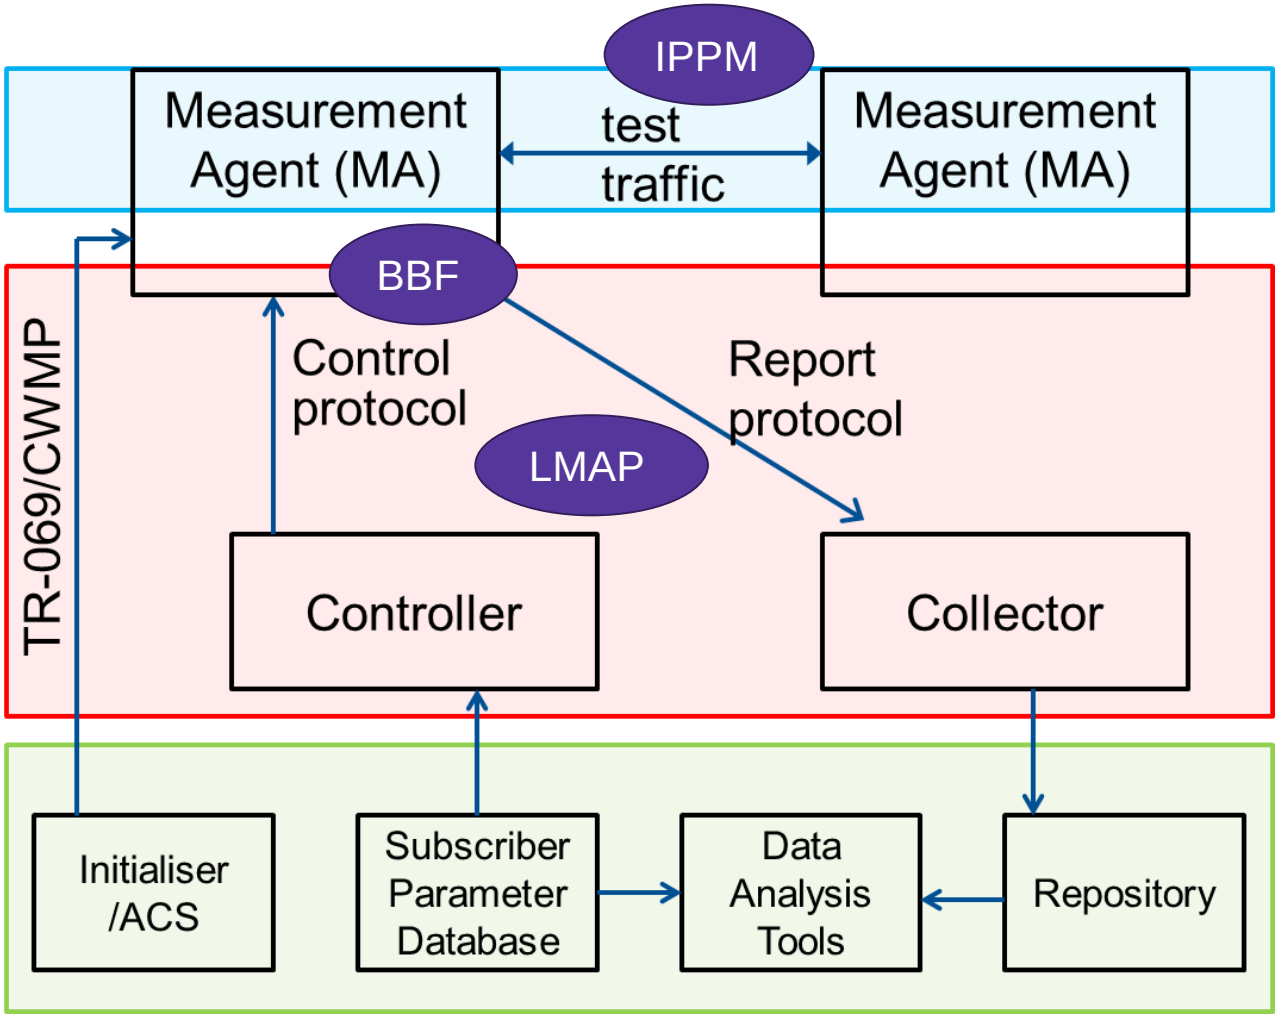

tests

IETF & BBF

BT's OAM

- Standards for large-scale, comparability and vendor interoperability
- Standard open about how results used, analysed, shared
- Limited progress on common tests