Next-Generation Measurement Architecture Standardization and Outreach Group (NMASOG)

MISSION STATEMENT

The goal of the Next-Generation Measurement Architecture Standardization and Outreach Group (NMASOG) is to facilitate the standardization and broad deployment of a comprehensive Internet broadband performance measurement architecture, protocols, metrics, and data structures. It provides requirements and input to standardization bodies, and describes how the output of these external organizations can be best used to create a reliable, accurate and sustainable broad performance measurement infrastructure.

NMASOG is initially focused on serving the needs of US residential broadband consumers, the FCC and carriers, but aims to make its recommendations broadly applicable beyond those groups, such as for enterprise networks and to evaluate performance beyond the "last mile". The components of this effort may also allow customers of broadband services to evaluate their own network performance and help with diagnosing performance-related problems.

The NMASOG addresses statistical and operational features of end to end broadband performance measurement, integrating and profiling standards and practices developed in external standards bodies.

NMASOG builds on the work of the Federal Communication Commission's (FCC) *Measuring Broadband America* (MBA) program and seeks to leverage other industry, academic and regulatory efforts. NMASOG will provide support for the development of concepts for measurement of performance and the promotion of the appropriate protocol and metric standards development in organizations such as the IETF, Broadband Forum, and other Internet and broadband technology groups. In addition, NMASOG will facilitate development of best practices and standards for operational practices, statistical models, measurement data processing, and presentation of measurement data, all of which are crucial to making broadband performance measurement relevant to consumers and industry.

NMASOG will also serve as a forum for other important discussion on topics that fall outside the scope of technical standards development. It will, for example, promote improved broadband performance knowledge in academic research, industry best practices, regulatory policy, and in other areas valuable to consumers.

OVERVIEW OF REQUIREMENTS AND ELEMENTS

We stipulate as a goal that a large number, if not all, consumer edge devices, such as cable or DSL modems, will be equipped with measurement capabilities "out of the box". At any given time, only a small fraction of such *measurement-enabled edge devices* (MEED) will actively participate in performance measurements, or may only participate in some portion of such measurements. For example, a large number of MEED devices could collect network availability (outage) data, while only a small fraction may conduct on-going active performance measurements. Other devices may engage in shorter-term collection efforts, e.g., to help a provider diagnose network problems. Below, we refer to a *measurement panel* as the group of MEEDs that are currently part of a performance measurement effort, such as MBA. MEEDs may conduct measurements outside of a panel, as noted. The selection of nodes to serve on the panel is determined by the entity managing the measurement project; NMASOG is expected to provide general guidance on principles for panel size, composition and methodology for the conduct of specific tests..

The MEED needs a measurement partner that receives and generates test traffic, and provides echo functionality for round-trip latency measurements. We refer to this device as an *in-network measurement server* (INMS). These servers can be operated by the ISP or a third party; a larger ISP may have several such servers. A mechanism needs to be defined for the MEED to discover the appropriate measurement server, most likely through the server described next.

The *Collection Control Administration server* (CCAS) provides the MEED with a schedule of tests to perform, possibly based on its characteristics, and the network addresses of the in-network measurement server(s) to be used as counterparties. In addition, the CCAS furnishes the MEED with a server where the MEED can deposit the measurement results (measurement log). The MEED may need to authenticate to the CCAS. The in-network measurement servers report their availability to the CCAS. The standardization effort would yield the format of the test schedule, e.g., as an XML document.

The *Measurement Data Collection Server* (MDCS) receives the uploaded data logs from the MEED and makes them available to the entity managing the data collection effort. (We are currently not concerned with the elements that anonymize, filter or aggregate data; there may be several such entities.) The standardization effort is expected to create a common log format for the measurement data that contains, inter alia, the MEED identifier, the time of measurement, the metric and the measurement result.

In order to be able to compare measured and advertised performance, we need an information service that can authoritatively determine the connection parameters of the broadband connection, such as the service provider (ISP), geographic location, network technology, service class, rate caps, burst capability and provisioned connection speed. As part of the standardization effort, we need to agree on a common identifier, such as a node identifier, circuit identifier or public IP address, that allows the MDCS, for example, to look up the customer subscription data. We refer to this logical server as the *Customer Network Information Server* (CNIS). It would likely draw on the business process systems of the ISP to obtain this data. Again, the retrieval protocol and data format are subjects for standardization.

In summary, the following metrics, protocols and data formats will likely be considered for possible standardization:

- Metrics, such as throughput and latency, typically using existing protocols such as ICMP, UDP or TCP along with metric-specific test parameters;
- A description (document format) of the test schedule and the tests to be invoked;
- A protocol to retrieve the test schedule document from the CCAS;
- A data format for conveying test logs to the MDCS;
- A protocol for uploading these test logs to the MDCS;
- A data format for describing the subscription characteristics of the ISP customer;
- A protocol for retrieving these characteristics.

In addition, NMASOG is expected to identify a set of operational requirements to ensure repeatable and reliable network performance measurements, including

- Minimum system requirements for MEEDs, such as on-device permanent storage (disk or flash memory), CPU performance and RAM;
- Network locations of MEEDs on the subscriber premises, as well as the number and placement of in-network measurement servers;
- Appropriate panel sizes and composition by technology and provisioned connection speeds;
- Measurement schedules, such as the frequency and the metric-specific parameters of tests;
- Statistical data processing guidelines, such as outlier removal;

Identifying and successfully recruiting a set of measurement subscriber nodes to match a given testing plan is crucial to achieving a statistically meaningful survey of a panel's broadband performance. The architecture provides protocols, data structures, and best practices to define how potential measurement subscriber devices' demographic and broadband performance characteristics are described, identified, and

matched (recruited) to a proposed testing plan based on key measurement and demographic characteristics. Protocols provide for both mandatory and optional test features that enable and enforce client configured opt-out of certain proposed test plans. Schedules also accommodate other carrier and subscriber concerns such configuration of test features for data cap policies, time of day or other temporal restrictions, and content and traffic management concerns.

Measurement and reference server nodes communicate, authenticate to, and register with a Collection/Control/Administration (CCA) server node and report their testing capabilities. CCA nodes authenticate NSIS nodes and a panel and testing profile recorded at a CCA node is disseminated to NSIS nodes to query available subscriber nodes. In response the CCA node directs subscriber and measurement nodes to conduct the test plan and report results back to the CCA node. CCA node responds to queries and requests for data regarding a given test plan.

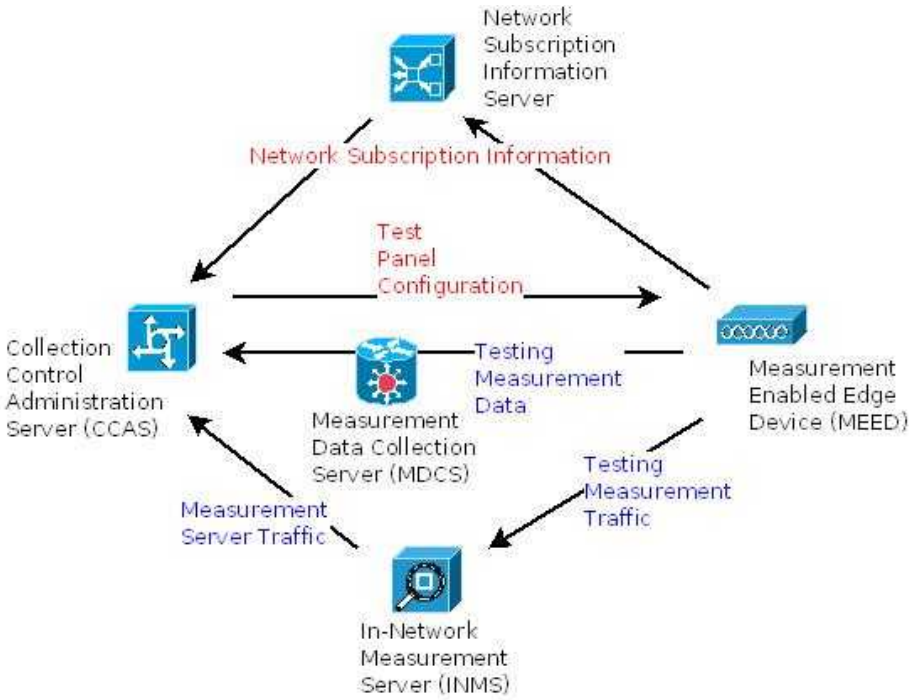

Figure 2

REFERENCE DESCRIPTION OF PERFORMANCE MEASUREMENT TECHNOLOGY

End to end consumer broadband performance depends widely on the diverse network hardware and software that traverse technology from a customer's own connected device, through broadband providers' last-mile networks, potentially any Internet connected network, to hosts operated by large Internet giants and other customer devices alike. NMASOG performance testing technology identify key end-to-end characteristics of common Internet usage and focus on 1) specific testing simulations of video streaming, web browsing and other common Internet use cases and application specific models; as well as 2) broader fundamental network performance characteristics such as DNS responsiveness, TCP data transfer rates, round trip latency, packet loss, or testable features. Standards and best practices define the parameters for implementing tests as well as underlying data structures and methods of describing and processing results of tests. Tests implement passive observational approaches to measuring broadband performance as well as active tests that rely on the measurement subscriber node generated artificial traffic.

INTERNET SERVICE PERFORMANCE TEST PARAMETERS

- Features
  - IPv6 support
  - DNSsec support
  - Port reachability
- Data rate
  - Sustained and burst
  - Upload and download
  - Sensitive to differing congestion algorithms
  - Capable of capturing TCP characteristics by including slow start
  - Total traffic and measured time period
- Latency
  - UDP echo round-trip time
  - ICMP echo round-trip time (“ping”)
  - Latency under load
- Reliability
  - Loss of connectivity (data link)
  - Packet loss
  - Delay variation (jitter)
  - Extreme delay
- Application emulation
  - DNS lookup delay
  - Video streaming (possibly adaptive)
  - Audio streaming
  - VoIP
  - Web metrics
    - Simulated web retrieval
    - Download time for top-N web sites