TECHNOLOGY TRADES IN IP-BASED TELEMETRY NETWORKS

**Joshua D. Kenney[1], Myron L. Moodie[1],
Gary L. Ragsdale, Ph.D.[1], and Thomas B. Grace[2]**

**(1) Southwest Research Institute®
San Antonio, Texas**

**(2) Naval Air Systems Command (NAVAIR)
Patuxent River, Maryland**

**joshua.kenney@swri.org, myron.moodie@swri.org,
gary.ragsdale@swri.org, thomas.grace@navy.mil**

ABSTRACT

The integrated Network Enhanced Telemetry (iNET) project established a test article standards working group to define open standards for network components and interfaces for test articles in the aeronautical test environment. Its aim is to utilize the growth of Internet technologies for telemetry networks and ensure interoperability among network devices. This paper describes the technology background and the current technology trades of an IP-based network paradigm used in producing standards for test article networks. Specifically, the paper will include descriptions of selected network technologies as applied to test data and time distribution within test articles.

KEY WORDS

iNET, Telemetry System, Network Protocols, Data Acquisition

IP-BASED TELEMETRY NETWORKS

For the proliferation of network-centric telemetry systems, a set of standard networking technologies is required to ensure interoperability between vendors of like components, economical sources of products, and flexibility of user applications. It is important to leverage standard technologies that are well established, with proven performance, continued future market support, and mature technologies with multiple competitive sources of integrated system components and constituent electronic devices (e.g., integrated circuits, wiring, and software).

The Internet Protocol (IP) suite is a common set of protocols used by the Internet and most existing commercial networks. This set of core technologies can be viewed as a set of layers, with each layer specifying a certain level of data transfer, using services from lower layer protocols, and providing well-defined services to upper layer protocols. The IP suite provides a core set of technologies necessary for standardized data transfer in network-centric telemetry systems.

IP-based telemetry networks use the IP network layer protocol to logically address hosts and route data packets throughout the infrastructure. IP is encapsulated by Ethernet, a data link layer protocol providing common hardware addressing and Media Access Control (MAC) for the transmission of data frames between network devices. Ethernet defines a number of wiring and signaling standards for physical layer transmission of 10 megabits-per-second (Mbps) to 10 gigabits-per-second (Gbps) over metallic or fiber optic cabling.

BASIC NETWORK CONNECTIVITY AND SERVICES

To achieve basic connectivity in IP-based telemetry networks, a certain set of protocols should be used at each layer to allow the simple transmission of network packets between connected devices. Standardizing this core pool of technologies provides a common set of technologies from which simple IP-based telemetry networks can be constructed. Utilizing these basic standards, more complex standards can be realized to address solutions with varying levels of interoperability and performance. The integrated Network Enhanced Telemetry (iNET) project is working to define a set of standards for network-centric standards, starting with technologies to establish basic network connectivity and services.

At the time of the writing of this paper, preliminary standards for basic network connectivity and services have been agreed upon by the iNET test article standards working group. These protocols are seen as the predominant technologies for network-centric systems and provide a foundation for future work in IP-based telemetry systems. The basic network connectivity technologies and their sources are listed in Table 1, and the common network service technologies and their sources are listed in Table 2.

Table 1. Basic Network Connectivity Technologies

Protocol Layer	Technology	Source
Physical	• Metallic media: o 10Base-T Ethernet o 100Base-TX Ethernet o 1000Base-T Ethernet • Fiber optic media: o 10Base-FL Ethernet o 100Base-FX Ethernet o 100Base-LX10 Ethernet o 1000Base-SX Ethernet o 1000Base-LX Ethernet • Ethernet auto-negotiation	IEEE 802.3-2005
Data Link	• Ethernet Medium Access Control (MAC)	IEEE 802.3-2005
	• Logical Link Control (LLC)	IEEE 802.2-1998
	• Transparent bridging • Rapid Spanning Tree Protocol (RSTP)	IEEE 802.1D-2004
Network	• Internet Protocol (IP) version 4	RFC 791
	• IP over Ethernet	RFC 894
	• IP over IEEE 802	RFC 1042
	• Internet Control Message Protocol (ICMP)	RFC 792
	• Internet Group Management Protocol version 3 (IGMPv3)	RFC 3376
	• Requirements for Internet Hosts	RFC 1122
	• Requirements for IP version 4 Routers	RFC 1812
Transport	• Transmission Control Protocol (TCP)	RFC 793
	• User Datagram Protocol (UDP)	RFC 768
	• Requirements for Internet Hosts	RFC 1122
	• Requirements for IP version 4 Routers	RFC 1812

Table 2. Common Network Service Technologies

Service	Technology	Source
Address Resolution and Host Configuration	• Address Resolution Protocol (ARP)	RFC 826
	• Reverse Address Resolution Protocol (RARP)	RFC 903
	• Internet Standard Subnetting Procedure	RFC 950
	• Bootstrap Protocol (BootP)	RFC 961
	• Dynamic Host Configuration Protocol (DHCP)	RFC 2131
	• Classless Inter-Domain Routing (CIDR)	RFC 4632 (BCP)
	• Requirements for Internet Hosts	RFC 1122
	• Requirements for IP version 4 Routers	RFC 1812
Name Services	• Domain Name System (DNS)	RFC 1034 RFC 1035
File Transfer	• File Transfer Protocol (FTP)	RFC 959

TIME SYNCHRONIZATION

In addition to the basic tier of networking technologies already covered, IP-based telemetry networks also require accurate time synchronization of distributed device clocks. A high-precision network time synchronization capability makes it possible for data sources to apply accurate time stamps to acquired data before transport on the network. This capability is critical for enabling time correlation of data at data sinks and simultaneous or scheduled sampling at data sources given the variable-but-constrained-latency characteristic of network-based data acquisition. Although several protocols exist to synchronize clocks over a shared network medium, only the Precision Time Protocol (PTP) allows the sub-microsecond accuracy necessary for most telemetry systems. PTP is recommended by the iNET Proposed Architecture and is defined in the Institute of Electrical and Electronics Engineers (IEEE) 1588-2002 standard. For an in-depth look at the application of IEEE 1588 in IP-based telemetry systems, see [1].

Test article and ground station networks should use PTP master clocks synchronized by a highly-accurate external time source, such as the Global Positioning System (GPS). Attached network devices, such as data acquisition units (DAUs), recorders, wireless controllers, and legacy adapters, must implement PTP slave clocks in order to maintain synchronization with the master clock. Depending on application requirements, design choices can be made for slave clocks, trading higher accuracy, higher cost hardware implementations against lower accuracy, lower cost software implementations. Boundary clocks or transparency should be supported in routers and switches in order to eliminate large latency variations that are typically induced in these types of devices.

Current implementations of PTP-synchronized networks often provide one pulse-per-second (1 PPS) external outputs of device clocks to allow verification of time signal lock using oscilloscopes. Another best practice is to allow slave clocks to run freely using the last known time in the absence of a master clock. Likewise, master clocks should be allowed to run freely using the last known time in the absence of an external time source.

TRANSFER OF LATENCY/THROUGHPUT CRITICAL (LTC) DATA

Thus far, we have covered the technologies needed to provide basic network connectivity for IP-based telemetry networks along with a high-precision network synchronization protocol that enables high-resolution time-stamping for data correlation. However, additional considerations need to be made when developing telemetry networks that support the latency and throughput guarantees of traditional Pulse Code Modulation (PCM) based telemetry systems. When designing synchronous-like data transfer over asynchronous network transport, some performance-enhanced networking mechanisms must be employed to ensure the timely transfer of latency-critical data and maintain sufficient bandwidth for throughput-critical data flows. Using existing network technologies, this section covers an approach for the delivery of latency/throughput critical (LTC) data using common schemes for the routing and prioritization of data flows sharing the network communication channel.

The reliable and efficient delivery of LTC data can be realized using IP multicast streams. With this approach, data sources send one or more data flows to specific multicast destination addresses. Data sinks subscribe to the multicast destination addresses of interest, and the data is delivered by the network to only subscribed devices. The Internet Group Management Protocol version 3 (IGMPv3) is the core technology for managing multicast group memberships, handling group subscription requests by IP hosts and instructing IP routers to deliver data marked with multicast destination addresses to only the IP hosts in the corresponding multicast group. Using multicast as the approach to routing reduces the overhead associated with routing each data flow as an individual unicast transmission and is more bandwidth-efficient than the simple broadcast transmission of all data flows.

This LTC data multicast approach uses the User Datagram Protocol (UDP) for end-to-end transport. As such, no transport layer acknowledgements are available for reliability, and data may be delivered out-of-order to data sink applications if routers are in the path of the LTC data. To alleviate these shortfalls, a couple of design decisions must be made. First of all, this approach assumes that deterministic network switches (or routers) are used to partition the network such that only two devices, i.e., the switch/router and the end device, are sharing a specific Ethernet link. Therefore, no Ethernet frame collisions occur and each pair of network devices can operate in full-duplex at full link speed without contention from other devices. Furthermore, queuing buffers in switches (or routers) must be chosen sufficiently large to handle the "burstiness" of the routed traffic. Even with this level of forwarding determinism, which is inherent to most modern implementations, aggregate throughputs on each link of telemetry networks must be planned to not exceed the available bandwidth. Using the anticipated send rates of each of the data flows and the pool of available multicast addresses, data flows must be assigned to multicast addresses in proportions that do not overflow available line rates and buffer queues. This planning can be done manually based on the anticipated network topology or automated in an application for producing setup files. In either case, metadata must be generated and shared between data sources and sinks for associating data flows with their multicast groups. Using the metadata as a key, data sinks can subscribe to only the multicast groups containing the data they desire. The concept of grouping data flows into multicast streams for routing is illustrated in Figure 1. Finally, application-layer sequence numbers must be used on multicast packets to allow reordering at end applications and to detect lost packets. Even in the rare case of a loss (1 in 10^9 in current implementations), sequence numbers allow end applications to detect specific missing packets for logging and error reporting.

In addition to the deterministic planning of data delivery in IP-based telemetry networks, Quality of Service (QoS) markings and policies can be applied to data flows to augment service quality guarantees for different data types. The assignment of QoS levels can even be carried out on a per-measurement basis, with data flows and multicast streams inheriting the marking of their most critical constituent data. Most modern switches and routers employ QoS schemes such as Diffserv (RFC 2474) to detect and prioritize QoS-marked data packets in the case of scarce network resources (transmission links, forwarding queues, input/output buffers, etc.) It is up to users to classify the importance of their data flows and add appropriate QoS markings at the data sources to aid in-network prioritization in the cases of contention.

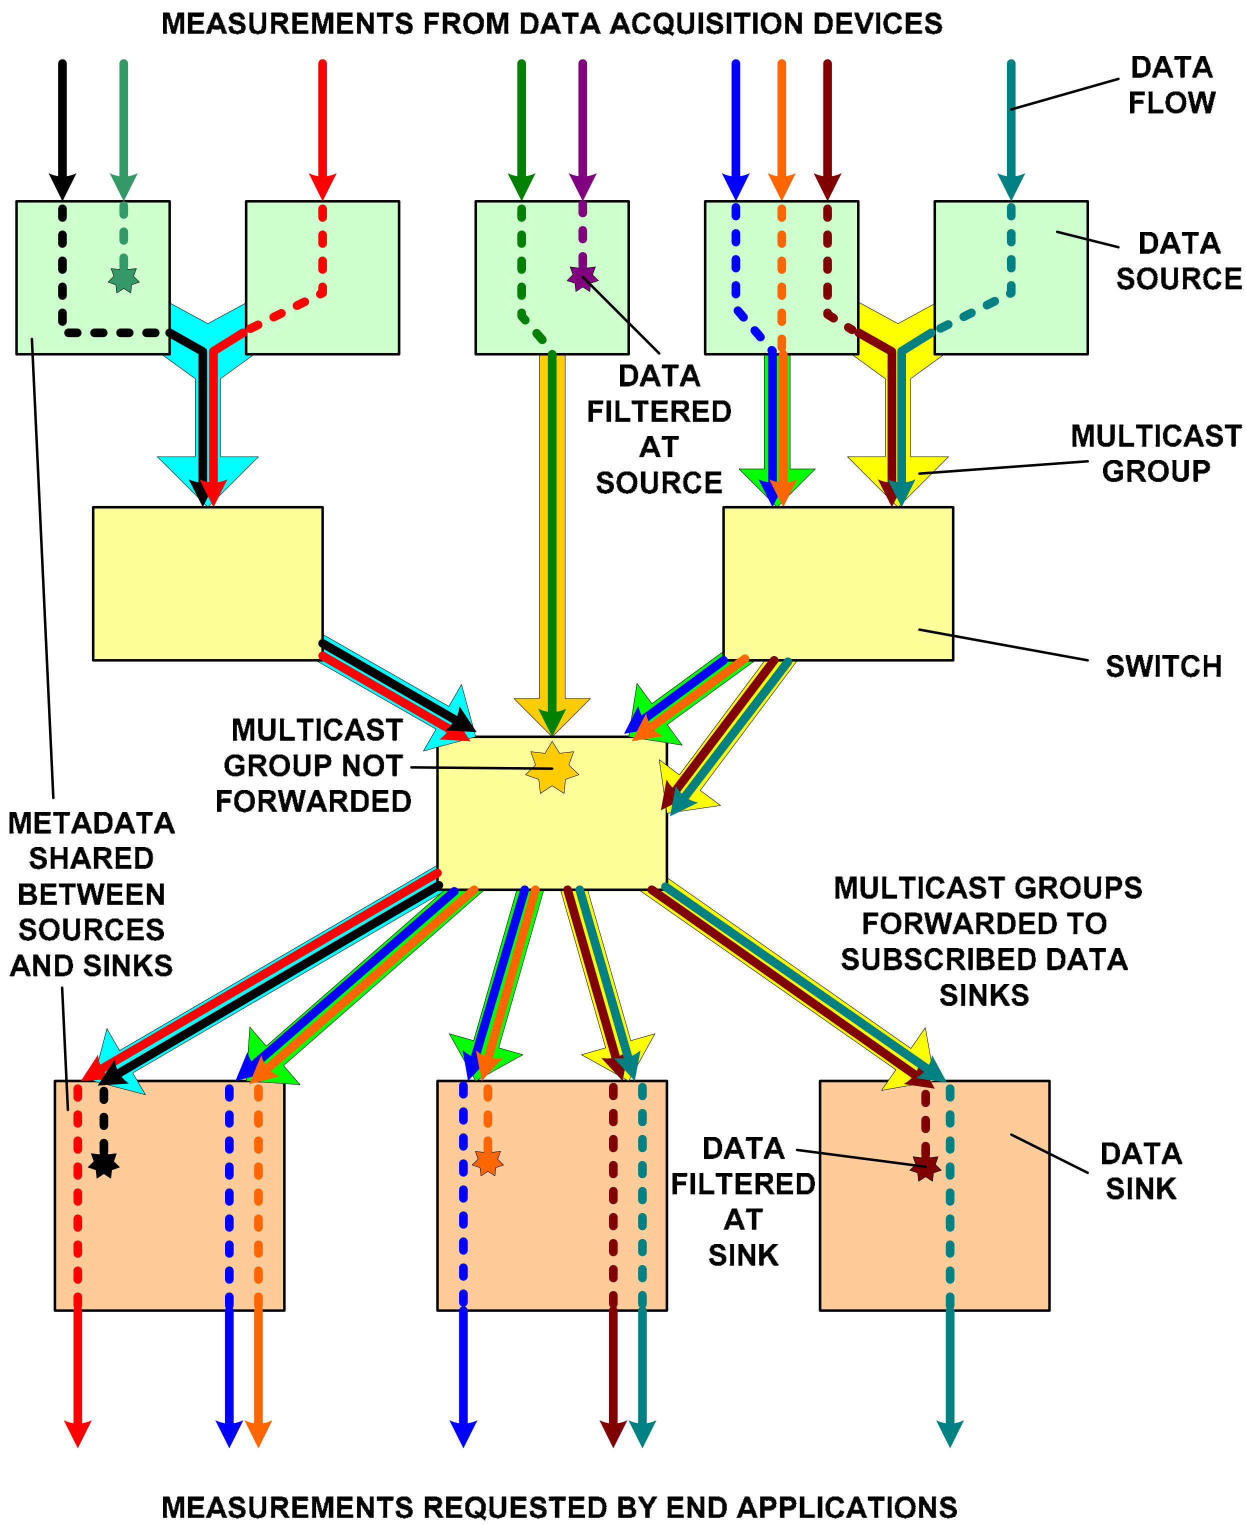

Figure 1. Multicast Routing and Data Flow Filtering

ENABLING APPLICATION LEVEL INTEROPERABILITY

Using the aforementioned techniques, an IP-based telemetry system can be realized to handle LTC data. However, we must also take a look at some mechanisms above the transport layer. First of all, applications must be aware of standardized methods of formatting messages before sending IP multicast packets to configured destination groups. We will also take a look at how filtering can be implemented at data sources and sinks to select data flows and measurements of interest.

From the approach presented thus far, we see that LTC data packets sent on the telemetry network need standard fields for time-stamps, application-layer sequence numbers, and message types (to differentiate LTC data from other data). Since standard IP multicast packet headers do not support these fields, a standard payload format needs to be determined. At the time of the writing of this paper, the iNET test article standards working group is discussing the standardization of custom payload formats to meet project requirements for telemetry network standards. The working group is also considering the adoption or adaptation of existing standard packet formats such as the Airbus IENA (Instrumentation d'Essais des Nouveaux Avions) format, the Inter-Range Instrumentation Group (IRIG) 106 formats, and/or several other proprietary formats that may become open to the general community. There is also an option to allow packet payload formats to be entirely described in metadata, resulting in flexibility in implementation choices at the cost of additional application complexity for handling a metadata configurable packet processing.

Assuming a particular implementation is using a standard packet payload format (or metadata configured packet payload format) that supports the fields needed for the transfer of LTC data, we will outline the task of filling a payload body with sampled data. To ensure the management of latency and throughput data flows, data sources can use a relatively simple approach to decide how to send sampled data. Limits of "maximum buffer size" and "maximum elapsed time" are configured at the data source and the minimum of the two variables is used to determine when to send the buffer of sampled data. Data sources send the contents of a sampled measurement buffer to a data flow when either the "maximum elapsed time" expires or the buffer reaches the "maximum buffer size." This results in the ability to "control" the sending of data flows based on a balance of latency and efficiency constraints. For example, setting a high "maximum buffer size" value (with an arbitrarily high "maximum elapsed time" value) is the best choice to achieve high network bandwidth utilization since packets are completely filled before sending and overhead is minimized. On the other hand, setting a low "maximum elapsed time" value (with an arbitrarily high "maximum buffer size" value) is the best choice to achieve low latency since packets are sent at a necessarily minimum delay, albeit with higher overhead caused by smaller packets. It should also be noted that sampled data from multiple acquisition measurements can be sent in the packet payloads of the same data flow. In this case, additional keys or common offsets should be used to partition the separate measurements for de-multiplexing at data sinks. The setup of these keys and offsets can be described in standardized metadata and distributed to devices for reference in data source generation and data sink processing.

Since we have discussed a unified approach for generation of data flows at source applications, and the transfer of those data flows across the network to subscribed sinks, we should now focus on how applications can filter the data. Metadata describing the data flows, whether it is explicitly used for configuration at run-time or implicitly designed into processing algorithms, provides the basis for filtering throughout IP-based telemetry networks. Using metadata described filtering methods for filling data flows with sampled measurements and assigning those data flows to multicast groups, data sources can filter which data is sent to the telemetry network. In the network, additional filtering occurs by the "pruning" of multicast trees (using IGMP "snooping") to send streams to only those data sinks that have subscribed to the multicast group. Data sinks use metadata to determine which data flows belong to which multicast groups and subscribe to the appropriate multicast groups based on the need of a particular part of the stream. Data sinks can further use the metadata to filter streams to select only the data flows, and even the specific measurements, of interest. The concept of filtering at the data sources, network fabric, and data sinks is illustrated in Figure 1.

SYSTEM INTERFACES

The focus of this paper thus far has been using IP-based telemetry networks to transfer data between distributed data sources and data sinks. This section aims to supplement the mechanisms covered thus far with considerations for common interfaces between the telemetry network and other parts of a typical telemetry system. This completes the picture of how the network-centric subsystem for data acquisition can be used as a basis for full IP-based telemetry systems.

We first must address the situation where test article networks are bridged with ground station networks (and other test article networks) using a radio frequency (RF) network. Although the RF network may use unique physical and link layer mechanisms for connecting multiple test article and ground station networks, the wireless network can effectively be viewed as either a link layer (layer 2) switch or a network layer (layer 3) router. If the RF network interface is designed as a layer 2 switch, connected test articles and ground stations are treated as a single IP network "cloud". In this case, the RF network is transparently treated as another switch in our network fabric, and our unified approach for multicasting from data sources to data sinks simply works across the entire system. That being said, implementers may desire additional control of the data sent across the RF link. In this situation, "proxy" network devices should be deployed on test article networks to down-select data flows traversing the RF network. Using the data delivery scheme described above, "proxy" devices effectively act as data sinks for the test article networks, filtering data according to the metadata based schemes outlined, and act as data sources for ground station and other test article networks, forwarding the selected data flows. If the RF interface is designed as a layer 3 router, multicast stream packets traversing the RF link must be handled by higher layer routing protocols such as the Border Gateway Protocol (BGP) and Protocol-Independent Multicast (PIM) methods. This approach naturally handles RF network bandwidth constraints by only forwarding multicast streams to test article and ground station networks based on static or dynamic routing tables. Currently, iNET test article and RF network standards working groups are jointly tasked with defining a unified approach to the transfer of data across this interface. This interface will also be greatly influenced by a separate working group formed specifically to handle iNET security considerations (i.e., encryption, authentication, authorization, etc.).

So far, producers and consumers of data on telemetry networks have been generically defined as "data sources" and "data sinks." These notional network pieces are effectively standardized "adapters" of devices using IP-based telemetry networks for the transfer of data. The iNET working group for test article standards has also been tasked with defining interfaces to several common telemetry subsystems. These interfaces will rely heavily upon the basic connectivity and performance-enhanced networking approaches covered above. We have already described how peripherals (DAUs, recorders, etc.) interface to the network as data sources and sinks, but additional considerations need to be made for standardized management and the adaptation of legacy (non-network-based) devices. Also, many network-based telemetry systems aim to use existing Serial Streaming Telemetry (SST) subsystems to carry a subset of acquired data. SST subsystem interfaces to test article networks, while logically acting as data sinks, may also need to filter data and will require standard methods for framing packet data into scheduled PCM streams. Finally, standards need to be defined for specialized data transfer and management through the use of Ground Service Equipment (GSE), which can act as both a data source and sink for system setup and testing.

The multifaceted mechanisms and interfaces for the transfer of data in IP-based telemetry networks require a holistic approach to system management. Network components such as switches, routers, and adapters require configuration and often real-time control. The network components should also respond to status queries and generate fault events when errors occur. Looking beyond network management, system management is necessary to oversee the behavior of network-connected devices such as peripherals, time sources, and interfaces to SST and RF network subsystems. As such, standard interfaces must be defined for telemetry network system management. This paper has also repeatedly pointed to the need for metadata to be shared throughout the network to enable application level interoperability. Coupling metadata with system management interfaces enables instrumentation interoperability through standard methods for device configuration and hardware setup. Multiple iNET working groups are cooperating to define these integrated test article, system management, and metadata standards.

CONCLUSION

The iNET test article standards working group is tasked with creating open standards for network components and interfaces of IP-based telemetry networks. This paper presents a snapshot of the working group's current standards approach, outlining considerations for basic network connectivity, network time synchronization, performance-enhanced data transfer, application interoperability support, and common system interfaces. The standards being developed are driven by functionality requirements of the telemetry network system architecture, while allowing for interoperability between equipment and open competition between vendors.

From the technology trades presented in this paper, it can be observed that test article standards first adopt existing technologies where appropriate. Although custom approaches are taken when no commercially available alternatives exist, constituent technologies with implementation experience serve as a basis for the standards. Furthermore, utilizing the layered approach of the Open Systems Interconnection (OSI) model enables modular adaptation and extensibility as new technologies are developed. The ultimate goal of the current approach is to provide common functionality and standard interfaces that enable flexible implementation and vendor-neutral interoperability for the proliferation of IP-based telemetry network systems.

REFERENCES

Please see the following resources for further information:

[1] Grim, Evan T., "Achieving High-Accuracy Time Distribution in Network-Centric Data Acquisition and Telemetry Systems with IEEE 1588," *Proceedings of the International Telemetry Conference*, 2006.

[2] *iNET Program Web Portal*. https://www.inetprogram.org/.

Referenced protocol standards can be found at websites for the Institute of Electrical and Electronics Engineers Standards Association (IEEE-SA) and the Internet Engineering Task Force (IETF) Request for Comments (RFC):

[3] http://www.ieee.org/web/standards/home/index.html

[4] http://www.ietf.org/rfc.html