

A Brief Overview of the IETF and the Internet Architecture: Past, Present and Future

David Meyer, Leslie Daigle,
and a cast of thousands
May, 2005

Agenda

- Introduction
 - High level introduction to the IETF's Mission and Scope
 - IETF Engineering and Architecture Philosophy
- Overview of the Internet Architecture (IA)
 - The idea here is to characterize the properties of the IA as Past, Present, or Future Internet Architectural principles¹, with an eye towards getting an idea of where we've been, where we are, and where we're going.
- A Brief Case Study – The World Wide Web
- A Few High-Level Conclusions

1. Or some combination of past, present, and/or future.

Acknowledgments

- This is the work of a cast of thousands
- In particular, the list of principles used here is taken in part from work done as part of the New Arch¹ project
 - <http://www.isi.edu/newarch>

1. Future-Generation Internet Architecture

From IETF Mission Statement (RFC3935, 2004)

“1. Mission Statement

The goal of the IETF is to make the Internet work better.

[snip]

The Internet: A large, heterogeneous collection of interconnected systems that can be used for communication of many different types between any interested parties connected to it. The term includes both the “core Internet” (ISP networks) and “edge Internet” (corporate and private networks, often connected via firewalls, NAT boxes, application layer gateways and similar devices). The Internet is a truly global network, reaching into just about every country in the world. The IETF community wants the Internet to succeed because we believe that the existence of the Internet, and its influence on economics, communication, and education, will help us to build a better human society.”

IETF Scope

“In attempting to resolve the question of the IETF's scope, perhaps the fairest balance is struck by this formulation: "protocols and practices for which secure and scalable implementations are expected to have wide deployment and interoperation on the Internet, or to form part of the infrastructure of the Internet.””

What does all of this mean?

- Well, first, the mission statement is
 - derived from experience
 - i.e., not a simple list of desiderata
- And protocol specifications document developments
 - Necessarily, that which is new or changed
 - Noticing that the architecture is often implicit
 - We'll talk more about this in a minute
 - Open and transparent participation is key

IETF Core Values

- Pragmatic Engineering
- Cooperation
- Best Technology
- Maximize user value from the network
- Includes experimental and evolutionary network

Engineering and Evolution

- Modularity
 - We build the tools for deploying the Internet and its systems
 - Others can build and deploy a wide range of variations on the theme: working systems using the architected extensibility of those tools
 - Further extension must feed back into new requirements: life cycle
- Diversity
 - In the history of the Internet, there never has been one perspective that was ubiquitously deployed
 - Engineering for change while focusing on stability

The IETF & Engineering

- One toolbox, one collaborative fabric; multiple creations
- But we also have a caretaker role
 - The IETF remains a venue for sharing decades of collected engineering experience and applying it to the incremental development of Internet infrastructure and protocols for the good of all.
 - The IETF does not seek to become a single locus of control or aspire to prescribe specific service metrics for network behavior at any level.
 - But we are on the alert for proposals which may benefit some subset of the network while penalizing the whole.

End-to-End Principle(s)

- Before diving into this one, a few comments:
 - First, the end-to-end principle is perhaps the most fundamental and least understood of the Internet's architectural principles...Think:
Nothing should be done in the network that can be efficiently done in an end-system.
A function that must be performed at a higher layer should not also be performed at a lower layer (without a good reason)¹.
- What does this mean?
 - Example: What can best be done in a end system?
 - Transport flow control (e.g., TCP)
 - Note that while TCP performance depends to on buffering in the network, the end-systems can (to large extent) measure and adapt to network buffering
 - Example: What can't be done in a host?
 - End-to-End QoS

1. Definition due to Noel Chiappa, http://users.exis.net/~jnc/tech/end_end.html

End-to-end Principle(s)

- Generality
 - Past: The network is built with no knowledge of, or support for, any specific application or class of application
 - Present: Must account for NAT, firewalls, mid-boxes, manageability
 - Future: Controlled transparency, e.g., Protocol Normalization
- Robustness
 - Past, Present, Future: The end nodes are responsible for communications functions that can be accomplished entirely by those (end) nodes
- Fate Sharing
 - Past: State that is specific to a particular data flow between communicating end nodes is maintained by those end nodes
 - Present, Future: Must account for NAT, firewalls, mid-boxes, manageability

Properties of the Internet Architecture

- Before diving into this, let's just note that perhaps the most fundamental property of the Internet Architecture is that **its semantics are intentionally loosely defined**. In particular:
 1. The IA contains *no careful definition of the end-to-end semantics* of data carriage
 2. An instantiation of the IA carries user data *transparently*
 3. While the IA originally had *no model of its own performance*, there is quite a bit of activity around the development of architectural tools to measure performance
 4. The IA has *simplicity* as a guiding principle
- With that background, let's take a look at some of the principles underlying the Internet Architecture (IA)

Principles of the IA

- **Simplicity**
 - Past, Present, Future: Things should be as simple as possible, but no simpler
- **Multiplexing**
 - Past, Present, Future: The current Internet uses variable length packets as the universal approach to multiplexing data streams.
 - Present, Future: Circuit Oriented Packet Switched (co-ps), e.g., MPLS
- **Transparency**
 - Past: In the absence of transmission errors, user data that is delivered to the intended receiver is delivered w/o modification (“principle of minimum intervention”)
 - Present, Future: Must account for middle boxes, DPI engines, caches,...

Principles of the IA

- Universal connectivity
 - Past, Present, Future: A host can send data directly to any other host except when intentionally prohibited
 - Present, Future: Continuing development of DoS prevention and mitigation tools (e.g., GTSM)
- Immediate Delivery
 - Past, Present, Future: Connectivity is continuous
 - Past, Present, Future: In the absence of failure or overload, data is delivered immediately

Principles of the IA

- Subnet Heterogeneity
 - Past, Present, Future: The Internet makes minimal assumptions about link layer functionality
- Common Bearer Service
 - Past, Present, Future: The Internet provides a connectionless, end-to-end service
 - Must account for co-ps mode (e.g., MPLS)
 - Past, Present, Future: The common bearer service provides at least the minimal common service, i.e., the best effort service
 - Present, Future: Diffserv, Diffserv Aware MPLS-TE, NSIS,...
 - Past: There is no access protocol for end-systems (hosts)
 - Present, Future: SIP based signaling, NSIS, RSVP

Principles of the IA

- Connectionless Network
 - Past: The inter-network packet delivery mechanism is connectionless
 - Present, Future: Explicit path/source routing (MPLS)
- Minimal Dependency
 - Past: A minimum of network services and services is required to support end-to-end communication
 - Present, Future: Must account for IMS, DPI engines, Services Control Planes, ...
 - Past, Present, Future: The host/network interface is symmetric and there is no specific network access protocol, so that two Internet hosts can communicate directly without an intermediary
 - Future: SIP-based signaling, NSIS....

Principles of the IA

- Global Addressing
 - Past, Present, Future: The Internet uses a single global address space to identify the network attachment point(s) of each node. Packet forwarding decisions are based on these addresses
 - Present/Future: IPv4 and IPv6, NAT (virtual addressing)
 - Past, Present: IP addresses are overloaded as end-system identifiers
 - Future: Active investigation (e.g., HIP, SHIM6)

Principles of the IA

- Regions
 - Past, Present, Future: The Internet supports administrative regions (routing domains) for routing and policy control.
- Mobility
 - Past, Present: The Internet is optimized for non-mobile operation and uses a special case mechanism for mobility
 - Future: Active investigation (e.g., HIP)

Principles of the IA

- **Protocol Layering**

- Past, Present, Future: The Internet protocols are defined using layered abstractions (realized using a last-on, first-off “stack” of protocol headers)
- Future: Active investigation (e.g., “Role-based” architectures)

- **Distributed Control**

- Past, Present, Future: No single points of failure in the control plane
- Future: “Control Servers” (e.g., PCE)

- **Global routing computation**

- Past, Present, Future: The Internet performs a hierarchical, globally consistent routing computation to support loop-free, hop-by-hop forwarding of packets based on the destination address
- Present, Future: RSVP-TE, PCE, MH-PW routing (MH-PW forms a new “layer network” in G.805 parlance)
- Note: In-band control plane provides fate-sharing

Principles of the IA

- Security
 - Past, Present: The Internet has no mechanisms to constrain hosts that offer excessive traffic. In particular, it contains no defense against DDoS attacks.
 - Past, Present: The Internet does not include an architectural approach to protect its own elements from attack.
 - Past: Link encryption is sufficient and efficient for Internet security
 - Present: IPSec, GTSM, Identity/Mobility, lots of active work (sbgp, ...)
 - Future: Active work in many IETF areas

Principles of the IA

- Network Resource Allocation
 - Past, Present, Future: A transport protocol in an end node is responsible for sensing congestion and slowing its transmission when appropriate
 - Past, Present, Future: Transport protocols should be no more aggressive than TCP in the presence of congestion
 - Past, Present, Future: The Internet contains sufficient buffering to allow a host to operate a congestion adaptation algorithm with a round-trip of control latency
 - Present, Future: Diffserv, Diffserv-aware MPLS-TE, NSIS, RSVP,...

Case study

- World Wide Web
 - Not invented within the IETF or any other standards organization
 - No proposal lead through a gated review process for broad-scale deployment
 - By contrast, it succeeded by deployment in an open environment that allowed incremental buy-in, achieved critical mass, and took off
- Engineering
 - Successive refinements of underlying HTTP protocol within the IETF
- Architecture
 - Generation of applications identifier architecture (URIs)

Conclusions (1)

- The IETF, as an organization
 - Aims to be as neutral as possible about business models
 - Engineers for the common, open Internet
 - Weighs impact of deployment drivers and hurdles
- IETF participants
 - Often necessarily have a specific business model behind their technical perspective (drivers or hurdles)
- Almost by definition, creative tension and tradeoffs will ensue

Conclusions (2)

- The Internet is evolving to support diverse and emerging requirements sets
 - Including many of the requirements that are beginning to be specified by projects such as the ITU's NGN
 - Note this kind of “evolvability” is a fundamental property of the IA (deriving from its “minimalist semantics”)
- So where are the architectural misalignments?
 - And where can we work together?
- This is the topic for this session