Future Internet is Trusted and by Ethernet

In search for a new networking paradigm

Raimo.Kantola@tkk.fi
Dept of Communications and Networking

Grand Challenges in Networking

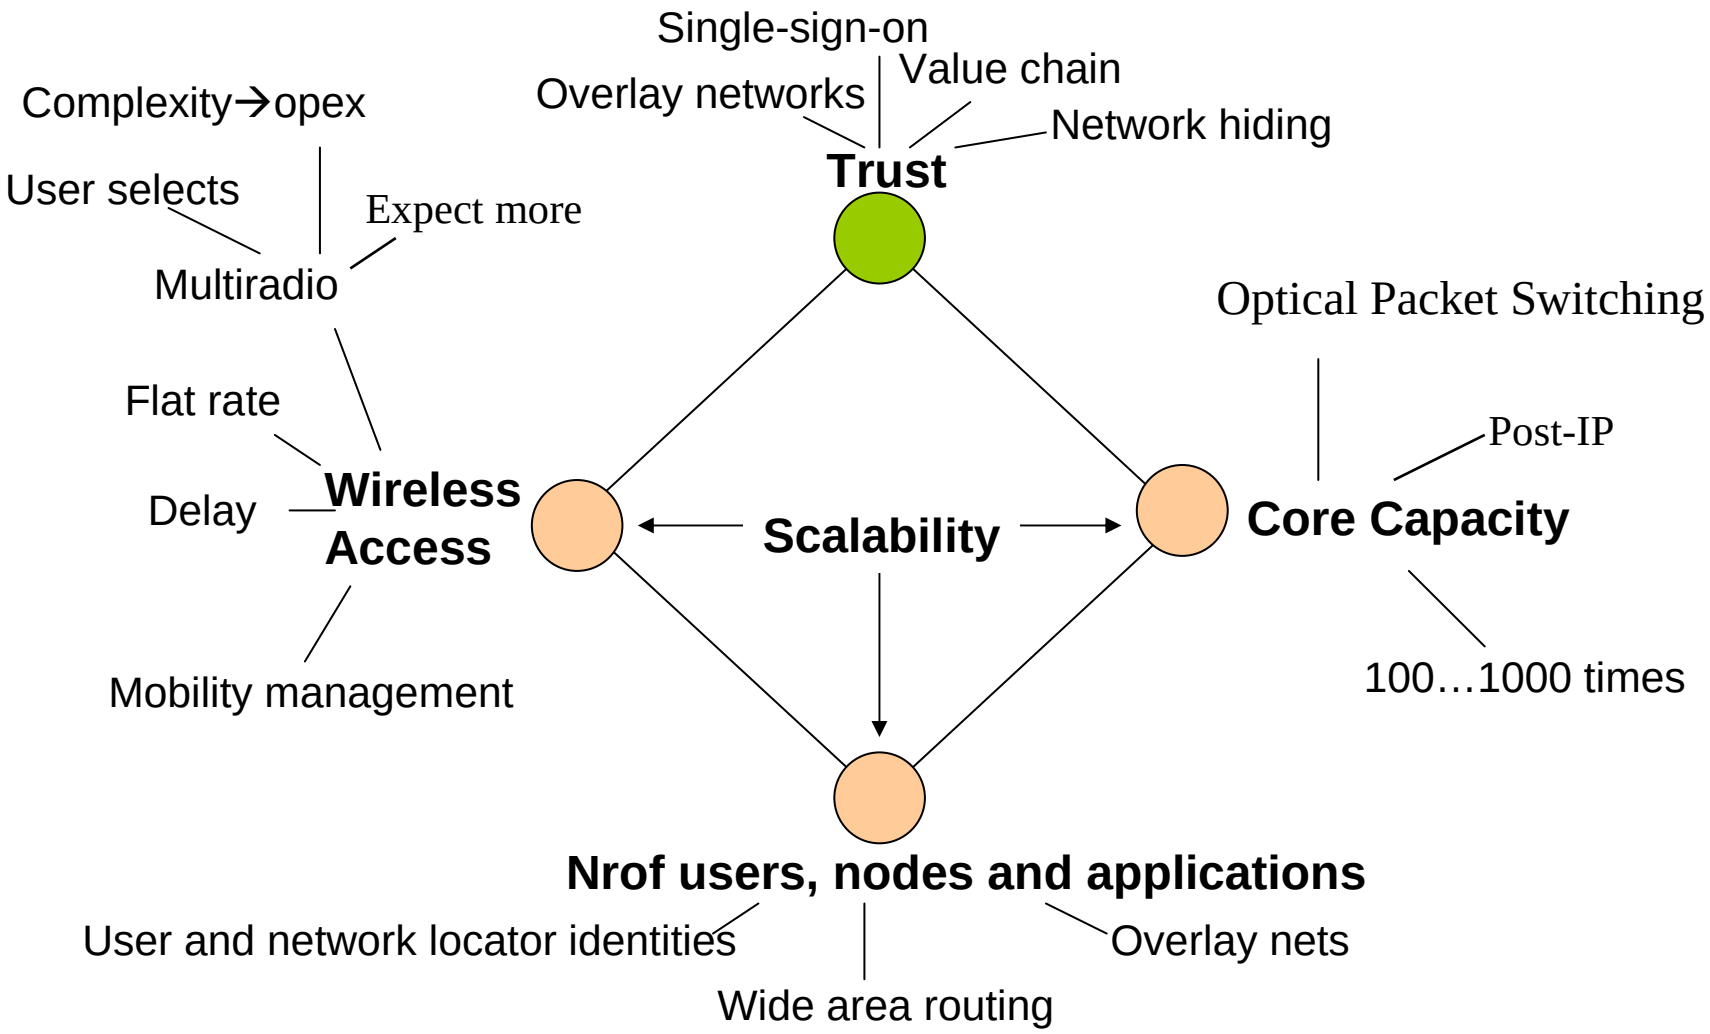

Agenda

- Current solutions
  - IP
  - Ethernet and its development
- Future solutions
  - by others
  - our solution: Internet by Ethernet
- Research issues and keys to success

Erosion of IP Principles

- Dave Clark, 1984: End to End Principle:

The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)
We call this line of reasoning against low-level function implementation the "end-to-end argument."

- Dave Clark, 2007: Trust-to-trust:

“The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at points where it can be trusted to perform its job properly.”
— David Clark, MIT Communications Futures Program, Bi-annual meeting, May 30-31, 2007, Philadelphia, PA.

IP over everything

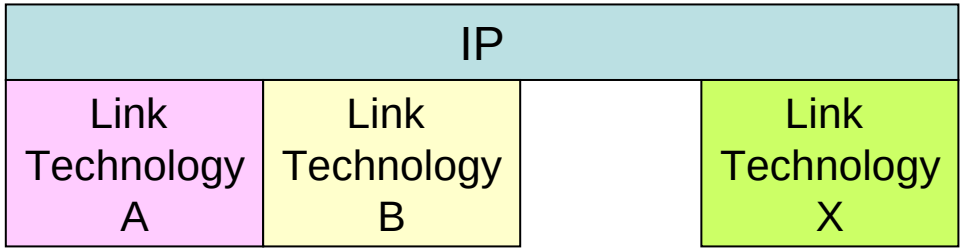

Original idea

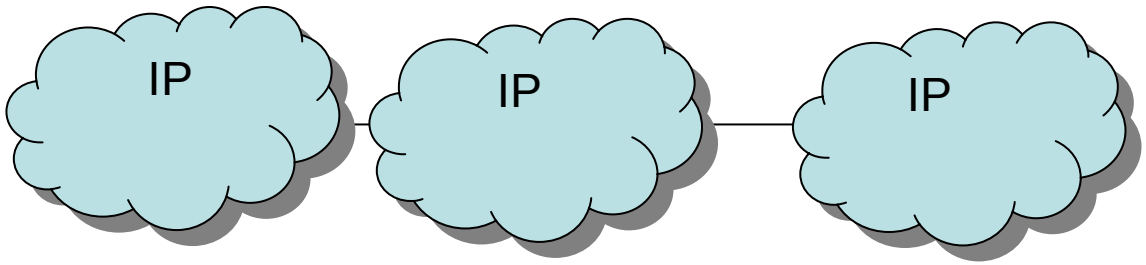

Reality today

- A lot of users have private addresses
- Users behind Firewalls
- Application Gateways between networks

- ”An IP connection” is made of legs belonging to networks that are hidden from each other.

IP Economics

- Flat rate service – economically efficient prices
- ISP margins from residential Internet services are low or non-existent
  - tends to grow uncontrollably creating a threath of forced investment and losses
- ISPs make their margin on Corporate connectivity services: VPNs etc.
  - close to half of traffic
- Best Effort and Flat rate facilitate misuse by unwanted traffic. Network pushes the cost of communication to the receiver.

What does this mean?

- Trust-to-Trust = We legitimise the misuse of IP-technology to meet legitimate customer needs in trust.
  - how to legitimize the breaking of IP over everything?
- Fundamentally, IP fails to support more than a single, low trust level. Customer and service needs in this respect vary greatly.
  - Lot's of add-on solutions have been developed.
  - The original IP network assumption that receiver wants to receive what sender sends is false → spam, malware.
- In the context of IP "Trust to trust" makes little sense. Rather, it is another "add-on".

What else is wrong with IP?

- IP itself does not support mobility
- Multihoming leads to fast routing table growth
- Visibility of networks to each other leads to long convergence or even instability of the routing system
- Middleboxes break many protocols → IETF spends a lot of effort in fixing the problems that emerge (NAT traversal etc…)
- ISPs have only few and obscure tools to map traffic onto their networks (MPLS, BGP).

Addressing – ”the IP matra” and the reality

- IP over everything – every host and server, every gadget has an address and therefore can be reached by anyone.
- This principle is contrary to reasonable real requirements: we want to control who can send packets to whom
  - DDOS protection of mobiles
  - our own home gadgets, control devices
  - corporate networks protected by NATs and Firewalls
  - hosts protected by Firewalls
  - use of Internet to control smart infrastructures (smartgrid etc)
- A huge addressing space (IPv6) is not a blessing – it is a problem.

NATs and NAT traversal break the E2E principle

- NATs were illegal in the eyes of IETF for a long time
  - save address space and protect local network from harm by hiding it
  - usually no incoming flow is admitted and only client-server applications work
  - For continuous reachability each application usually has to maintain its own mapping (connectivity state) in the NAT by some keep-alive signaling
- P2P started making fun of NATs → IETF reacted → first attempt at describing how NATs work failed → now second attempt
- IETF recommends UNSAF = UNilateral Self-Address Fixing = spy what the NAT is doing and adapt to it on application layer
  - Tools: STUN and TURN protocols
  - Solutions that use tools in a certain way: SIP Outbound and ICE (Interactive Connectivity Establishment)
  - **does not scale to mobile hosts** that want to be reachable e.g. for telephony or have a www-server or an active mp2p application
  - bog applications with code that has nothing to do with the actual task of the application

What is happening with Transport?

- ATM → does not scale → phase out
- SDH → scales up to 10 or 40 Gbit/s → not enough for future backbone links
- ISP requirement: Carrier Grade = ISP allows traffic from A to B, then it is transported. All other traffic is deleted.
  - IP is not carrier grade
  - MPLS tries to become carrier grade but MPLS is expensive (e.g. higher OPEX than SDH) – target is MPLS-TP
- Drivers for change are cost and sufficient capacity.

Ethernet development

- 10GE is shipping
- 100GE is on drawing boards, expected on markets in 2010/11
- Many new wireless access variants are emerging in the ”802.x” family
- Provider Backbone Transport (PBT) and Provider Backbone Bridging (PBB) are trying to become native packet based carrier grade transport solutions for network operators.
  - Connection oriented: route tables populated by Network management system
  - New variants of ”MPLS” used to support the creation of pipes and Traffic Engineering.

From bit stream transmission to packet transport

- Packet transport is more flexible in resource allocation than bit stream transmission, can easier make use of statistical multiplexing under heavy load and can easily adapt to a situation when there is nothing to send.
- In SDH and other bit stream transmission, one has to send and receive a constant speed bit stream irrespective of the user traffic, cost is constant → constant power consumption.
- Flexibility has been the driving system requirement for layer 3 and above and led to move from circuit switching to packet switching. Now it seems that flexibility will drive the change from circuit transmission to packet transport.

What is needed?

- ICT SHOK Future Internet Mission Statement:

  *Enhance the Internet technology and ecology as a platform for **innovation** while providing strong **governance over** the use of the network resources and information in such a way that especially **mobile use** of the network and its services will be natively supported.*

- IP over everything → encapsulation/ decapsulation on admin boundaries
- Transparent network → network as a black box
- Routing + DNS → Routing scalability + Switching for control + (trusted) Directory Services

Information networking

- Network segmentation reduces the Innovation potential supported by the original Internet model
- Users and companies coming up with solutions:
  - Peer-to-Peer
  - Data Oriented Network Architecture (ICSI)
  - Distributed web (HIIT)
  - Publish and Subscribe (P. Nikander, LME)
- Question is: on top of IP or without IP

Top-down or bottom-up?

Communication over Trust Domains

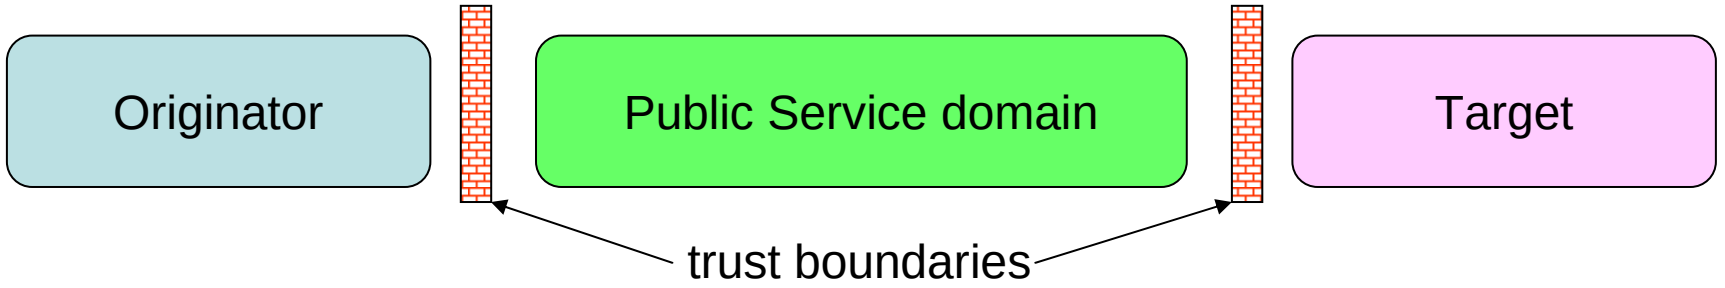

- Each trust domain has its own addressing = originator and target networks have private addressing
- ***None of the domains shall release any address information to other domains***
- For crossing a trust boundary, a packet shall present originator-id and target-id
  - a domain has ID Service for translating its IDs to addresses
  - processing at trust boundary is subject to policy

Communication Path is a Chain of Trust Domains

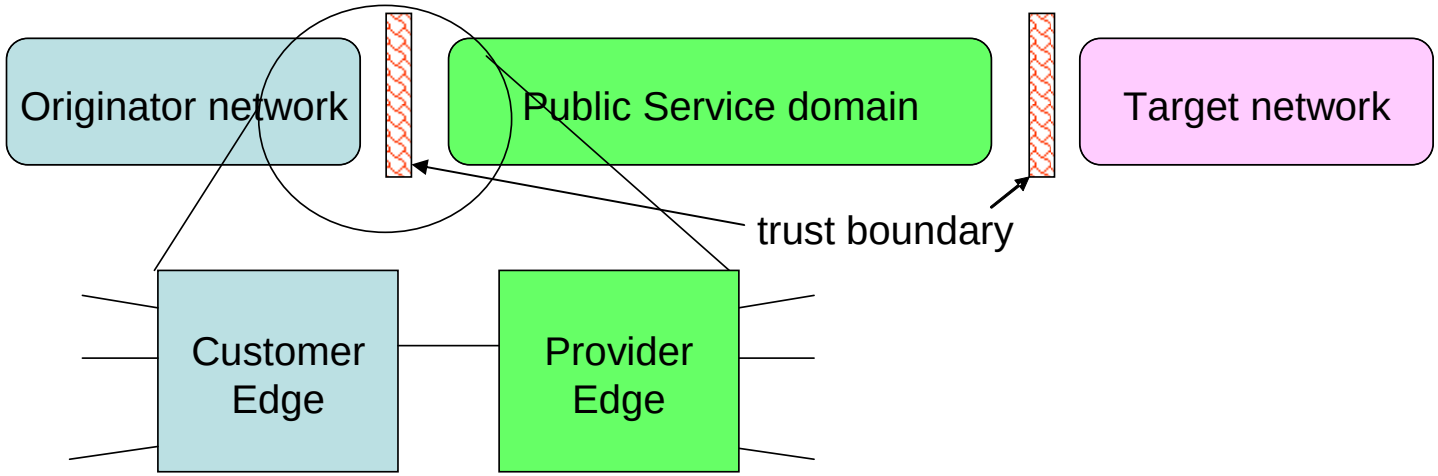

Trust boundary has Provider Edge and Customer Edge

+ policy processing has ***connection state*** → ”Customer Edge Switching”

+ state is managed by implicit signaling = signaling is embedded in the usual message pattern of ”DNS query – response – [message to target – response]*”.

+ Also Provider Edge will have ***connection state*** for hiding public service domain addresses from customer/user networks.

What kind of IDs?

- Globally unique and deterministic
  - high opex, operators do not want to maintain any new schemas, none exist (FQDNs are names not IDs)
- Random (sufficiently unique in any reasonable edge node)
  - feasibility: birthday paradox
  - managed by the home network
- Locally significant, static or dynamic
  - could be allocated by DHCP and carried in IPv4 address fields
  - Managed by the visited network

Birthday paradox: 1M users served by a Customer Edge Switch

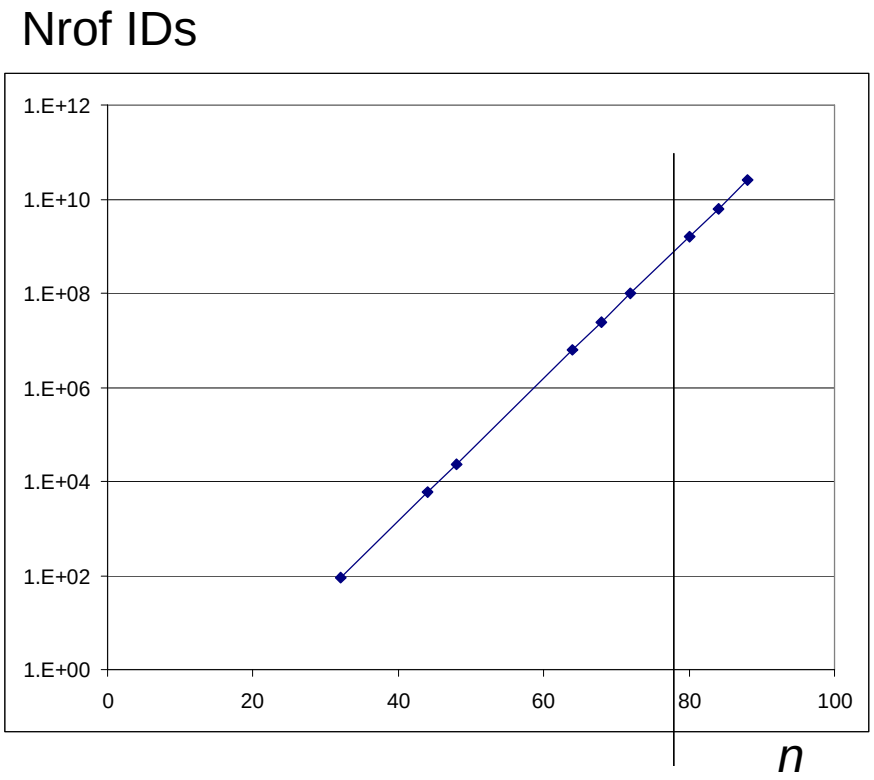

- BP = what is the probability that in a room with N people two have the same birthday
- Can be turned to: Given probability of clash *p*, how many IDs can come to a device if random ID length is *n*
- $p = 10^{-6}$
- ID dependent filtering can reduce *p* by another factor of 10^6

⇨ Reasonable Edge Switches can be built with IDs that have at least 60+ bits

Can Trust boundaries be chained (more than 2 on the path)?

- Answer: No and there is not need to.
- Why?
  - communication path has 2 ends.
  - does not scale to short flows that are typical of data traffic
  - would be equivalent to replacing IP by ATM that was supposed to be a switched broadband solution – did not work
  - a directory service can store addresses of PE for names of hosts and services but how could it save information about the whole path from PE to PE. If it tried, the DS would become responsible for routing.

Three Tier Program for Trusted Internet

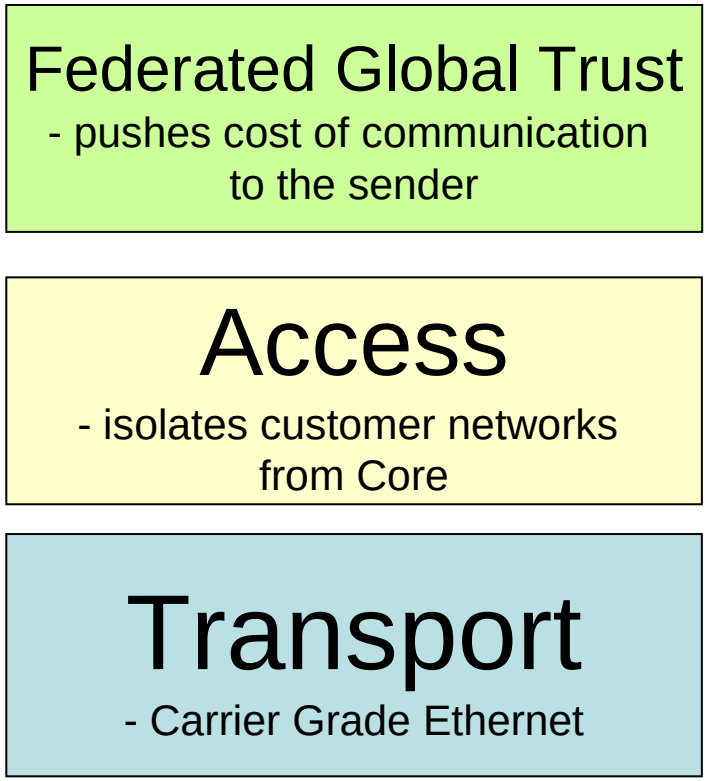

- Each tier can progress independent of the others
- The war against unwanted traffic can not be won by defense only
  → Global Trust System
- Access: key is supporting battery powered devices and mobility management
- Alternatives for the Core: CGE and IP/MPLS

Global Trust

- Target is to push cost of communication from the receiver to the sender
- How to do it?
  - access: stateful trust processing on trust boundaries
  - collect incident info, aggregate it, calculate trust index
  - rate providers and customers based on trust index
- Will this work?
  - can malicious senders attack the trust system?
  - theorem: Bysantine generals → in principle, if vast majority of users are ”good guys” they can isolate and win over the ”bad guys”.

Current IP/RE comparison

Today

IP + MPLS + VPN

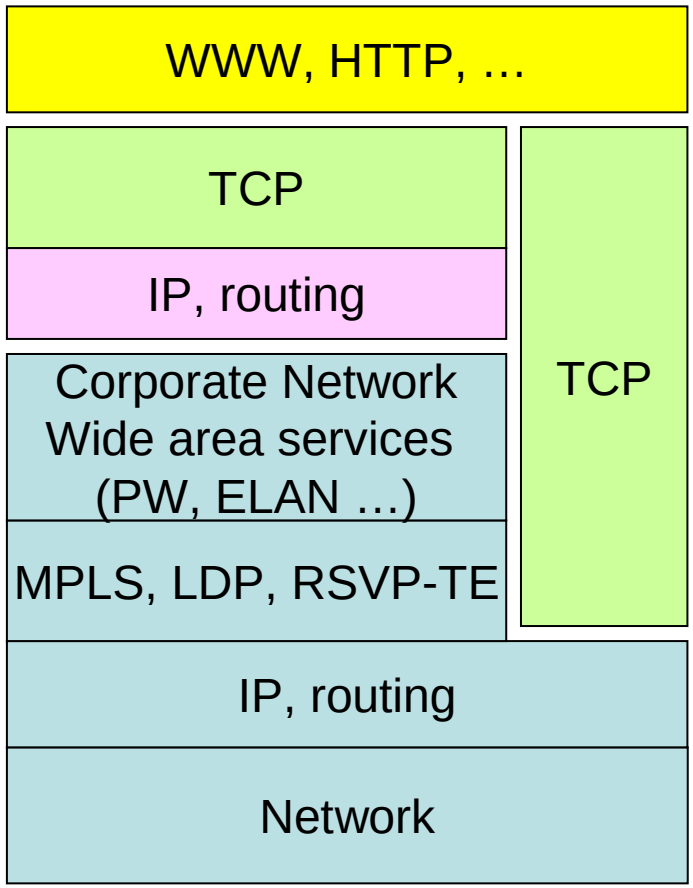

Future

RE2EE

Routed End-to-End Ethernet

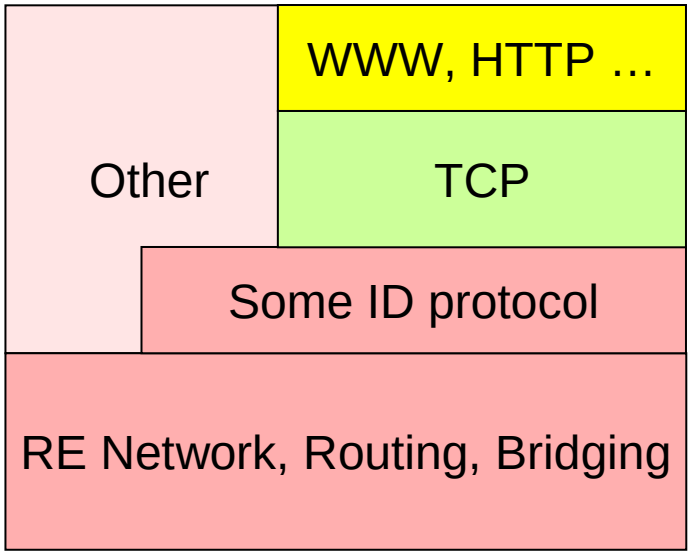

Most suitable stack: IP over MAC-in-MAC

Principles

Internet (based on IP)

- End-to-End Principle
  - E.g. DNS is a service among others
- IP over everything
- Network unconcious about users
- Dynamic routing
  - IP address has dual semantics
  - Support for a single naming/addressing scheme (IPv4 addresses **or** IPv6 addresses)
  - Multihoming visible to routing
- Data Plane and Control Plane not separated
  - All nodes visible to each other
- Mobility – at best some sort of add on.
- VPN support is an add-on (MPLS, IPSEC, etc)

Internet by Ethernet

- Trust to Trust Principle
  - Addressing and address resolution are an integral part of the network
  - A network does not publish its addresses to other networks owned by other admins
- Ethernet Everywhere
- Dynamic routing + dynamic address resolution + switching on the edge
  - Identities and locator addesses are clearly separted
  - Can simultaneously support many addressing schemas (IPv4, IPv6, NSAP, E.164 …)
  - Multihoming is an address resolution and matter of edge switching, does not impact routing
- Control Plane clearly separated from data plane → more robust design
  - Network not visible to users
- Mobility management implemented uniformly with other forwarding features
- Integrated VPN support, several parallel models for managing connectivity meeting different trust needs

Research Issues

- Addressing, identities and naming
  - Translations between the three, different ID schemas, how best to turn IP addresses into private addresses
- Control Plane
  - Routing (ISIS and other)
  - Service discovery by hosts
  - Address resolution
  - TE (we need to manage capacity allocated to different services and VPNs)
  - Mobility Management
  - IP networking over RE and to/from RE,
  - Switched Ethernet compatibility
- How to tackle security and unwanted traffic: packet access control and trust management
- Testing and deployment scenarios

Internet by Ethernet Benefits (1)

- Lower cost than IP based on wider economies of scale, lower stack and uniformity of design
- Better integration of routing with L2
  - Faster convergence, improved scalability
- Mobility
  - a uniform design with other MAC in MAC features
  - Tight integration with network attach/detach that are L2 features
- Robustness
  - Separation of Control and Data plane
  - Network is invisible to users, Core routing does not react to client network state changes
  - Neat support for multihoming with a combination of routing, edge switching and address resolution
- Uniform approach to services
  - → service tag → VPNs, TE, multitopology routing etc…
  - → Services available to mobile and non-mobile users

Internet by Ethernet Benefits (2)

- Separation of identities from locator semantics
  - → E.g. Identities can be per service or a set of services.
  - → Network assured identities can be required per service
  - → Anonymous identities can be explicitly protected by the network or a service (if one can trust the network or the service) (e.g browsing without being personally tied to having visited a site) without compromising efficiency.
- Differentiation by trust level
  - policy controlled trust at the edge on the trust boundary
    - ”Trusted” (e.g for VPNs, banking etc)
    - ”Operator assured” ID with AAA (e.g. similar to GSM)
    - Normal for BE services (e.g based on DHT) for web 2.0 etc.
    - mechanisms for pushing cost of communication to the sender
- For operator: divide and conquer – use one infra to support many services in a managed way.

Reality check

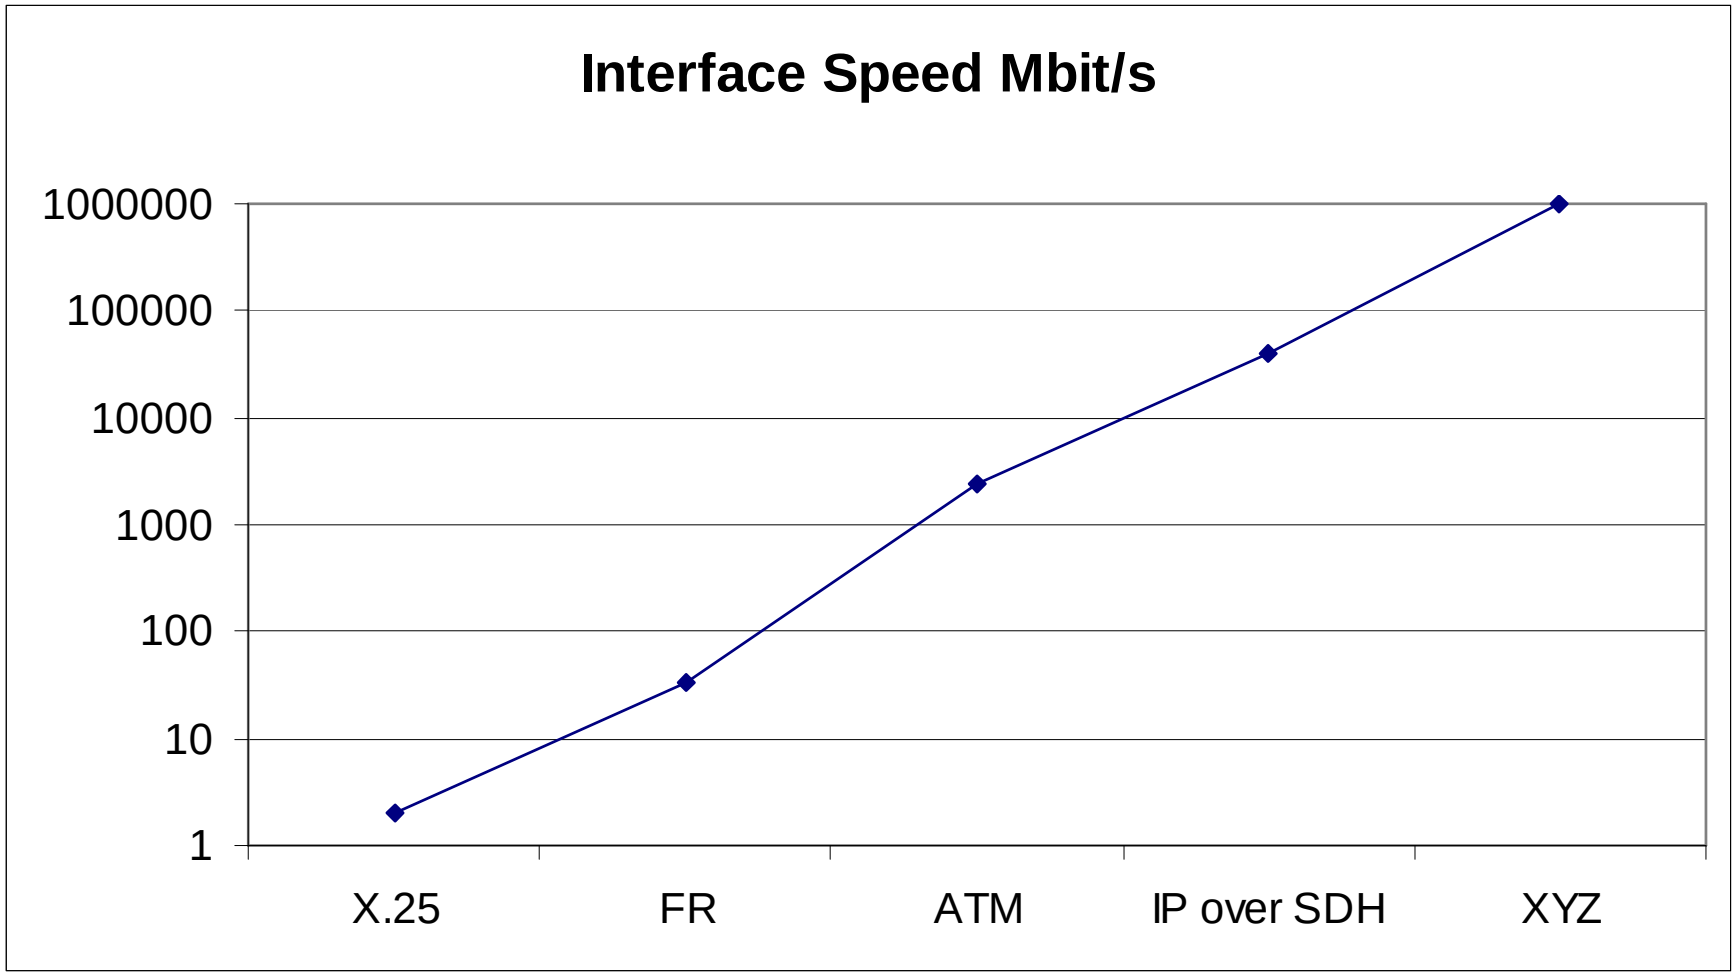

Summary

- Key terms
  - Routed End-to-End Ethernet (RE2EE)
  - Post IP, Internet by Ethernet
  - Uniform network technology providing unprecedented economies of scale and low cost trust-worthy BB services to the masses
- Global communication based on ***Globally unique names, local addresses and local identities***.
- Does not depend on upgrading all hosts connected to the Internet, rather emerges gradually from Metro Ethernet
  - Deployment in hosts and different networks is independent
  - 3 independent deployment areas: Core, Access and Global Trust
- Allows wide area networking without IP for a new service e.g. with network assured identities.

Are the benefits enough?

- The presented concept solves all well known IP network problems.
- Keys to success
  - Deployment scenarios: any player who invests can benefit and thus justify investment is key. The scenarios need to be developed.
  - Cost: Ethernet transport per 10Mbit/s should be less expensive than SDH
  - Adoption for mobile access: if this takes place the ball will start rolling.
- We have implemented CGE a'la MPLS-TP without IP that is suitable for intra and inter-carrier Core in ETNA.

Conclusion

- Search for a new networking paradigm is on
  - IP delivers the opposite to what are the stakeholder needs in trust
  - Information networking
  - Routed End-to-End Ethernet
- Ethernet replacing SDH for Core Transport
  - can we leverage this move to replace also IP?
  - Interleave Ethernet Carrier Grade pipes with pipes supporting dynamic routing and connectionless traffic
  - Is there an incremental path to enhance Ethernet?
- It is time to think beyond IP.