3.Diversifying the Internet

- Why diversify?
- Sample metanets
- Network architecture issues
- Substrate routers and metarouters

Ossification of the Internet

- Why diversify?
- Sample metanets
- Network architecture
- Substrate & metarouters

- As networks get hard to change
  - »investment in physical infrastructure
  - »many competing stake-holders
  - »even modest changes (e.g. IPv6) are difficult to deploy
- Internet ossification is a growing concern
  - »the internet, while useful, is far from perfect
  - »if network architecture can't change, forced to work outside the architecture
- Not unique to networking, but more severe
  - »dominant players in other domains can be slow to change
  - »but, competition allows new entrants to innovate and forces incumbents to respond, preventing ossification
- Networking less open to effective competition
  - »integrating virtualization into architecture can change that

GENI and Internet Diversification

- GENI advocates virtualization for experimentation
  - »expect new Internet architecture to emerge that will ultimately replace current architecture
  - »not clear how this migration will occur
  - »what about next time we need to change Internet architecture?
- Make virtualization part of Internet architecture
  - »enable multiple architectures to co-exist within Internet
  - »new networks can be added at anytime
  - »individuals or organizations can develop/deploy network
    - no need for universal agreement
- GENI can be vehicle for diversifying the Internet
  - »if GENI hosts several novel metanets with compelling apps, expect ISPs to offer commercial version of GENI
  - »critical issue is maintaining openness to new metanets

Diversifying the Net

- Diversification allows many *diverse metanetworks* to co-exist within shared *substrate*
  - » different metanets may differ radically
    - different protocols and service models
- More than just shared links
  - » diversified routers serve as hosts for multiple *metarouters*
  - » network processors, FPGAs deliver flexible processing
  - » new industry standards enable assembly of diversified routers from open, board level subsystems
- Substrate provides resource provisioning
  - » assign processing resources to metarouters
  - » connect metarouters via metalinks
  - » support dynamic connection (and reconnection) of hosts/devices to metanets

Diversified Internet Terminology

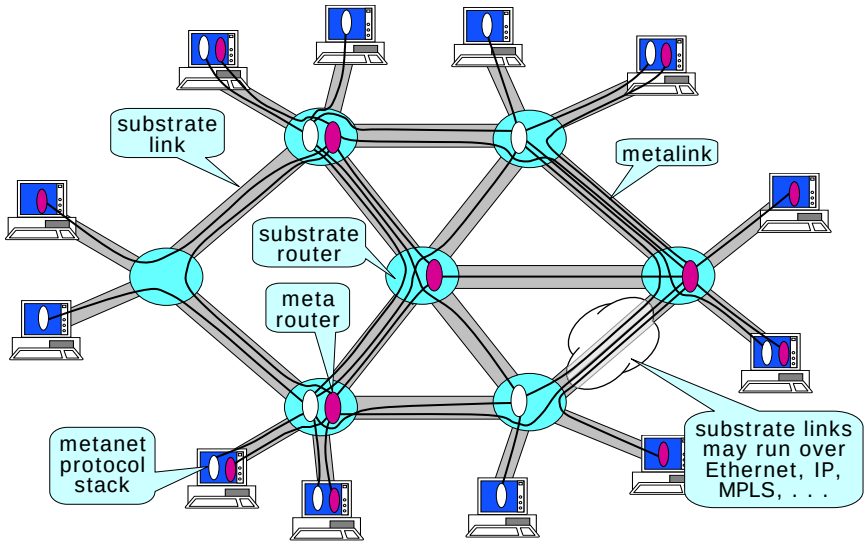

Moving the Waist

- Diversification shifts the “waist of the hourglass”.
  - » diversification layer lies just below traditional IP layer
  - » operates over multiple layer 1, 2 network technologies
  - » allows multiple end-to-end packet delivery systems
  - » focus is on providing resources used by metanetworks

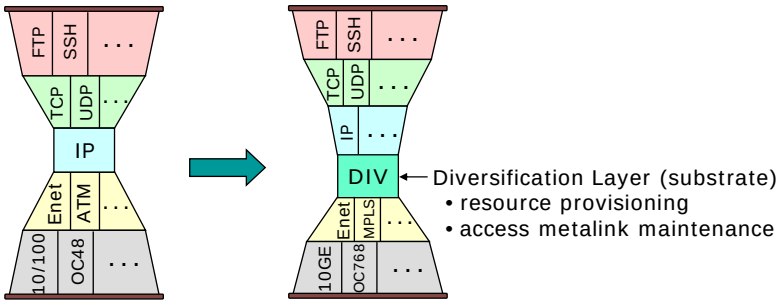

Uses of Diversified Internet

- Why diversify?
- Sample metanets
- Network architecture
- Substrate & metarouters

- Complete unfinished business.
  - »IPv6,IPSEC,IntServ,multicast,AQM, . . .
- Wide-area enterprise networks.
  - »network-embedded deployment of special capabilities
  - »better security, interaction with customers, suppliers, ...
- Enable introduction of new network architectures.
  - »enabling user control over traffic – e.g., i3, TVA, ...
  - »new services – reliable multicast, bulk data transfer, ...
  - »application-oriented nets – e.g., multimedia conferencing
- New opportunities for providers/vendors
  - »metanet providers innovate to create new markets
  - »substrate providers compete by providing best infrastructure and services to metanets
  - »equipment vendors differentiate by providing most capable platform and best metarouter design services

Scalable Distributed Simulations

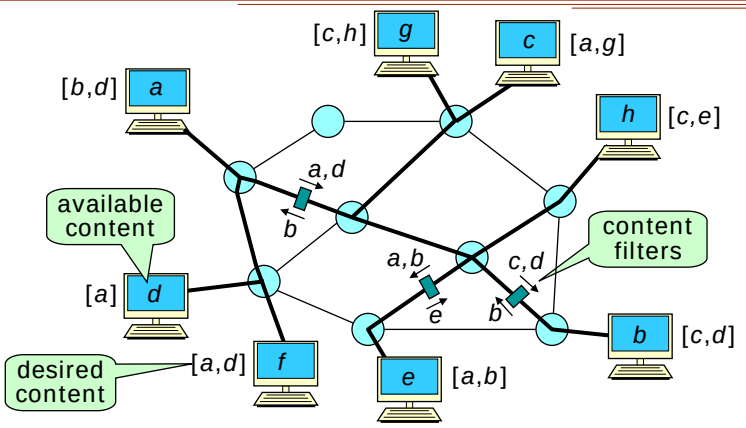

- Informed update propagation for dist. simulation.
  - »many-to-many multicast channel
  - »interest filters limit update propagation
- From Zabele, et. al. [2001].

DoS Attack Mitigation

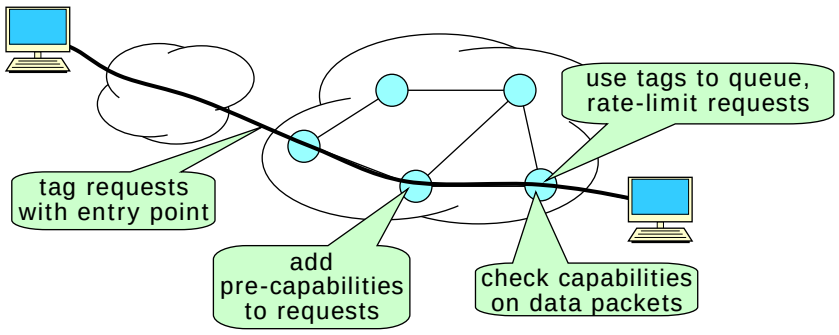

- Must get receiver's permission before sending to it.
- Routers block transmissions lacking *capability*.
  - »routers create piece of capability that only they can check
- Request rate limited by queueing based on entry.
- From Yang, Weatherall and Anderson [2005].

Video Narrowcasting

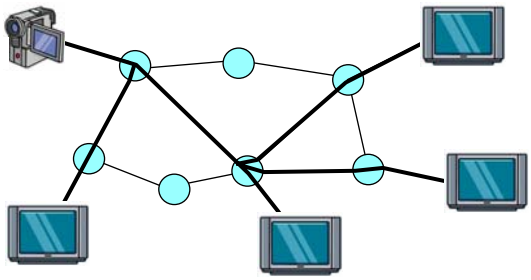

- Allow end-users to distribute real-time video
  - »high school sports, birthday parties, conference, . . .
- Requires network support for multicast and QoS
  - »allow "owner" to control access
  - »video format translation in network
- Directory services to allow users to find programs
- Pay for by advertisement – ad insertion by routers

A Bulk Data Transfer Metanet

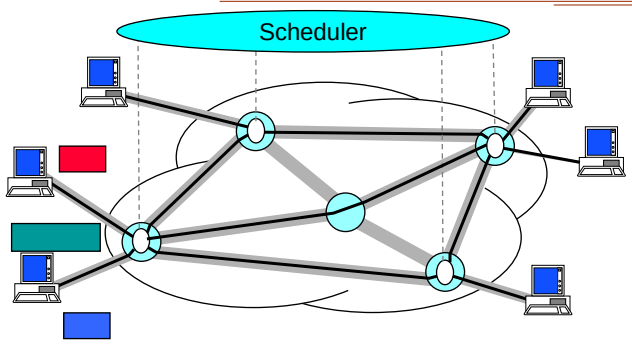

- Users request bulk data transfers in advance.
  - »intended for 10 GB to 100 TB transfers
  - »specify availability time and target delivery time
  - »network schedules *pickup*, *transfer* and *delivery*
  - »scheduled reservation ensures contention-free transfer
- May add temporary metalinks for large transfers.

Multiple Substrate Domains

- Why diversify?
- Sample metanets
- Network architecture
- Substrate & metarouters

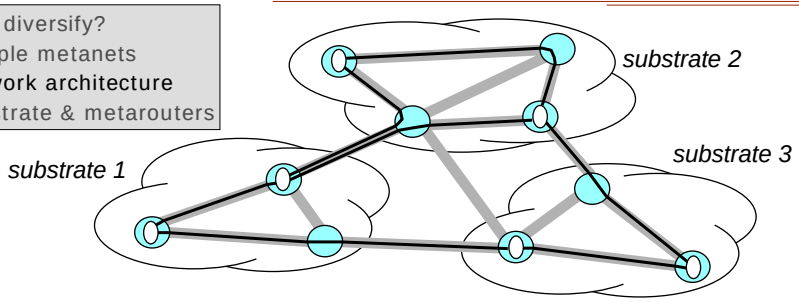

- Substrate naturally divides along geographic and organizational boundaries
  - »different ISPs will operate own substrates
  - »private organizations, individuals will operate substrates
- Metanetworks need not divide similarly
  - »single metanet may span the globe
  - »metanets can determine their handling of multiple domains

Metanet Configuration

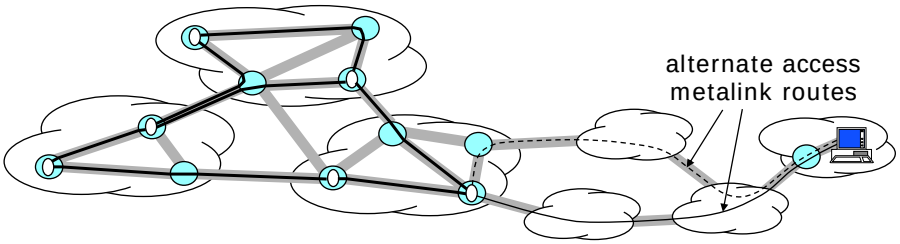

- Key issues
  - » backbone provisioning – metarouter processing, metalinks
    - sufficient for expected traffic, but responsive to unexpected
  - » dynamic access metalink configuration
    - users may connect from anywhere, anytime
- Impact of multiple substrates
  - » interact with substrate providers to obtain resources
  - » varying levels of trust, acceptable admin. complexity

Diversifying the Hosts

- Objective – enable hosts to use multiple metanets
  - »minimize barriers – at most, download/install software
  - »ensure isolation among meta-stacks
- Issues
  - »networking software currently built into kernel, no isolation
  - »standard OS mechanisms inhibit multiple protocol stacks
  - »hard to get good performance from user-level net stacks
    - many competing processes, coarse-grained time-sharing
- Alternate approaches
  - »secure, loadable kernel modules
    - language based, hw protection (segmented VM), hybrids
  - »user-level network stacks
    - application library for packet formatting, daemon for control
    - plus, configurable kernel mechanisms for queueing, filtering

Diversifying LANs

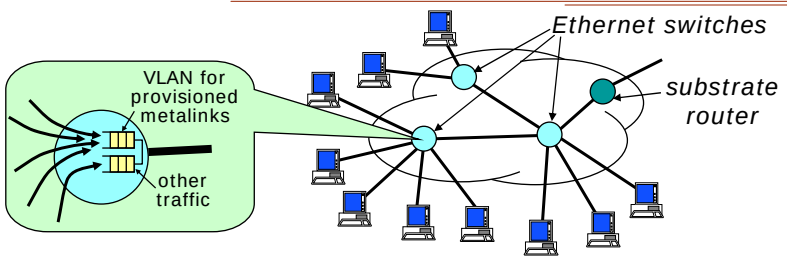

- Context
  - »almost all wired hosts connect via Ethernet
    - even 802.11 hosts are effectively Ethernet-connected
  - »to extend substrate to end systems, need to use Ethernet
    - VLANs and QoS queueing provide required mechanisms
- Substrate router regulates bandwidth usage
  - »directly controls downstream traffic
  - »allocates bandwidth to hosts, monitors for compliance

Security in Diversified Internet

- Alternate approaches
  - »security is fundamental – applies to all metanets
  - »allow metanets & substrates to take different approaches
    - including little or no security measures, support for anonymity
- Cost of securing substrate is weak argument
  - »encryption already commonplace in wireless LANs
- Technical issues
  - »what's built into substrate is hard to change and will ossify
  - »secure substrate reduces effort required of metanets
  - »metanets may require substrate assistance against DoS
- Policy issues
  - »governments' desire to protect (or control) citizens
  - »privacy concerns and availability of technology to ensure it
  - »business' competing needs for security and ease-of-use

Mobility in Diversified Internet

- Role of substrate vs. role of metanet
  - »allow metanets to handle endpoint connection process and mobility in different ways
  - »minimize substrate role to postpone substrate ossification
- Issues
  - »may be no metarouter near host connection point
  - »substrate may have little or no information about host
  - »metarouter may know little about substrate
- High level approach
  - »connecting device requests connection to desired metanet from a nearby substrate router
  - »substrate sends request to metanet, giving device location
  - »metanet contacts substrates to configure metalink
  - »as device moves, metanet reconfigures metalink as needed

Identity and Authentication

- Alternative approaches
  - »all devices have authenticatable identity
  - »devices may be anonymous, substrates and/or metanets determine identity requirements
- Pros and cons similar to security issues
  - »perhaps stronger case for identification in substrate
- Issues for weak authentication approach
  - »may be multiple authentication authorities
    - preferably small number (<< # of metanets or substrates)
  - »may refuse communication from unauthenticated hosts
  - »substrates/metanets *may* require authentication of connecting hosts
  - »some metanets may support anonymous communication

Substrate Addressing

- Substrate addresses provide location information to assist routing of metalinks (not packets)
- Hierarchical addresses merits/drawbacks
  - » hierarchy facilitates scaling
  - » but, networks are only loosely hierarchical
  - » connection to geographical proximity very rough
    - corporate nets with common prefix may span globe
    - ISPs with different prefixes may cover common area
  - » mobility is making location a central concern
  - » geography also important for metanet configuration
- Alternate approach include location with address
  - » (domain, location, host or service)
  - » domain+host unique within given location
  - » multi-scale specifications to match specificity needs

Control Communication

- Metanet-to-substrate
  - » for configuring metarouters, metalinks
- Substrate-to-substrate
  - » for configuring metalink segments
- Communication with general services
  - » e.g. authentication services
- Host-to-metanet
  - » to establish initial connection
- Support addressing of services, not just hosts
  - » allows metanet to contact substrate configuration service at a well-known address, based on substrate domain id
- Should be resistant to DoS attacks
- Can be implemented using a special metanet

Substrate Routers & Metarouters

- Why diversify?
- Sample metanets
- Network architecture
- Substrate & metarouters

- Substrate router must host multiple metarouters
  - » provides generic processing resources
    - compute blades, network processor blades, FPGA subsystems
  - » scale total performance by adding processing blades
  - » connect blades through flexible switching infrastructure
  - » provide isolation among metarouters
  - » diversified line card approach vs. processing pool
- Building metarouters
  - » like building physical routers using similar components
  - » key issue is making metarouter construction easier, without sacrificing (much) flexibility or performance
  - » enable use of multiple types of processing resources

Diversified Line Card Approach

- Similar to conventional router architecture
  - »line cards connected by a switch fabric
  - »traffic makes a single pass through the switch fabric
- Requires *fine-grained diversification*
  - »line cards must support multiple *meta line cards*
  - »requires intra-component resource sharing and traffic isolation

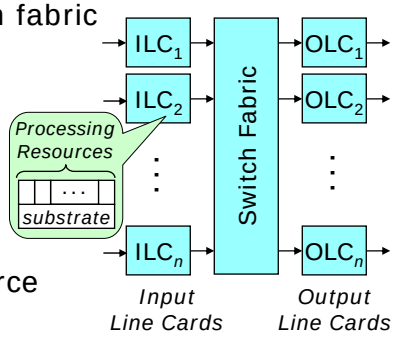

- Mismatch for current devices
  - »multi-core NPs lack memory protection mechanisms
  - »lack of tools and protection mechanisms for independent, partial FPGA designs
- Hard to vary ratio of processing to IO

Processing Pool Architecture

- Processing Engines (PE)
  - » used by metarouters
  - » variety of types
    - general purpose (GPE)
    - network processor (NPE)
    - FPGA (FPE)
  - » raw mode or cooked
    - cooked mode allows shared use
- Line Cards and Switch Fabric implement substrate.
  - » move packets to/from PEs
  - » isolation among metarouters
  - » flexible traffic flow in metarouters
    - PEs within MRs manage internal congestion