A Connection Oriented Internet Architecture for Restricting Reachability

Sneha Kumar Kasera

School of Computing
University of Utah

Introduction

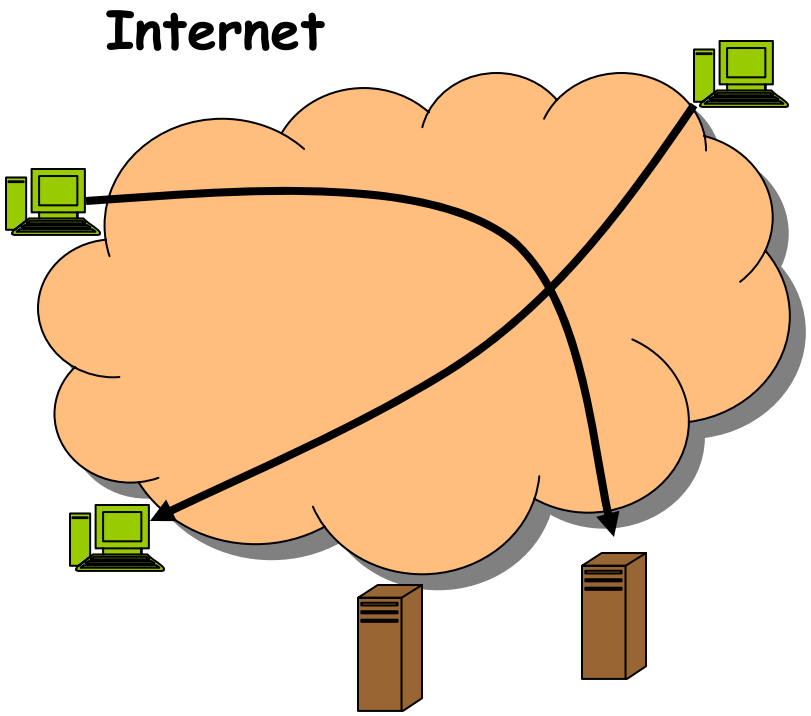

- any node with public IP address can be reached
- **reachability** can result in
  - vulnerability to port scans, digital pests
  - flooding, *slow-poison* attacks
- firewalls, filters, secure bug-free end systems can only help so much

if a node can be reached, it will be reached

Our Proposal

- new connection-oriented architecture to restrict reachability
- client nodes do not have IP addresses, clients known by user addresses, e.g., kasera@cs.utah.edu
- basic model: servers have well-known IP addresses
- clients send signaling messages, hop-by-hop, to set up connection paths (like virtual circuits) to servers

Our Proposal (contd.)

- local identifiers assigned to connections at routers/switches, used for forwarding packets
- soft connection state at routers, expires unless refreshed (often)
- client reachable by server only during connection, cannot be reached once connection state expires

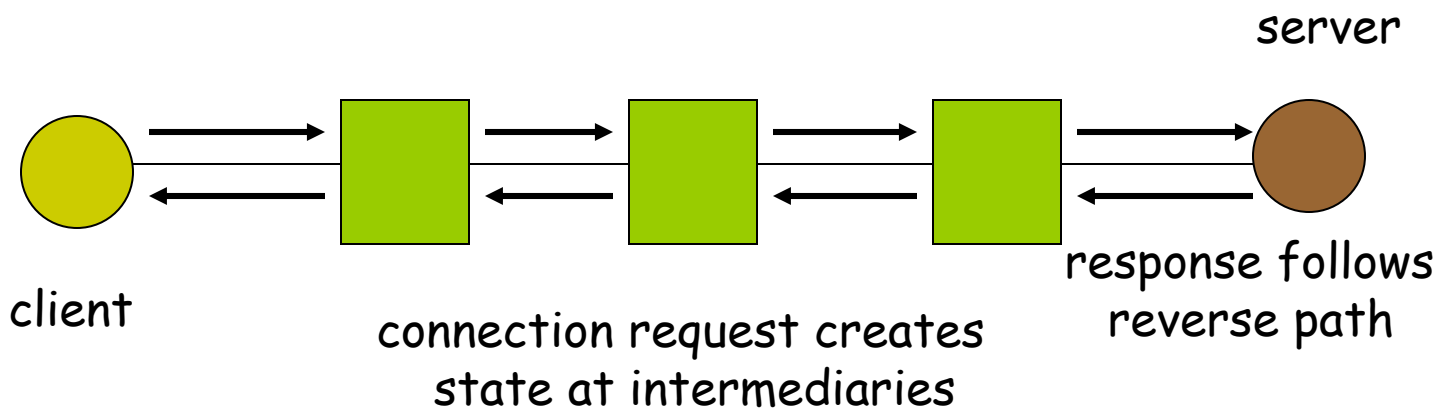

Securing Servers

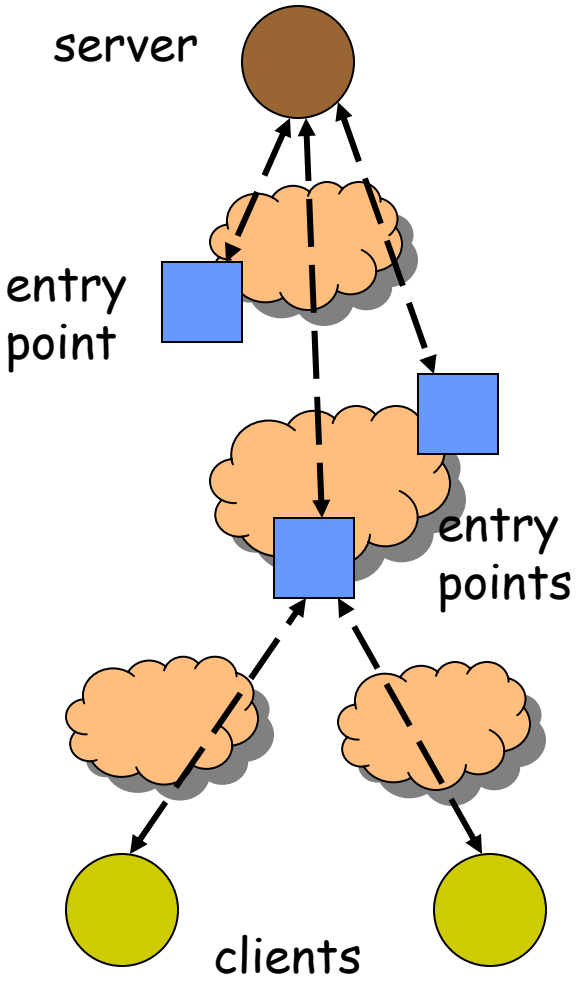

- enhanced model: servers do not have IP address, reached through well-known entry points
- clients connect to entry points
- servers set up connections to entry points
- entry points can be located anywhere, possibly in client domains

Peer-to-peer Communication

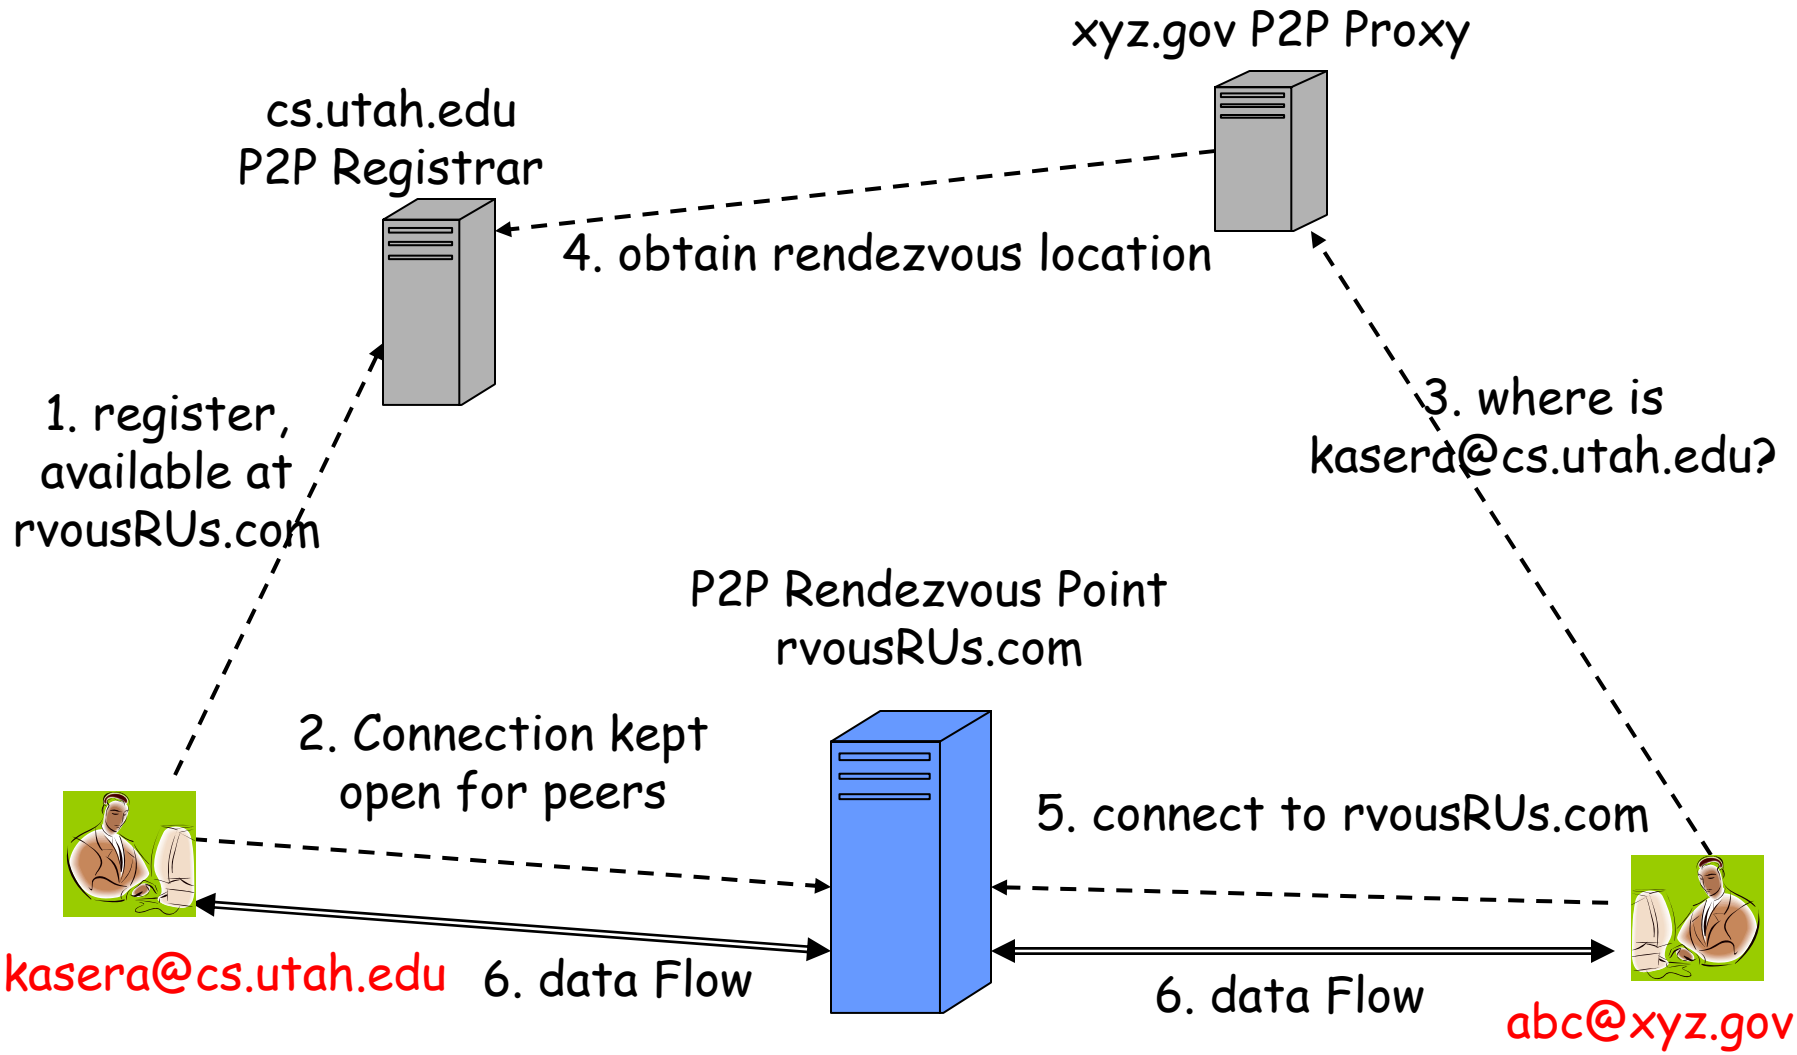

Architecture Benefits

- clients do not have public locators, cannot be reached when not active
- servers do not have public locators, can only be reached through entry points
- even entry points cannot reach servers when connection state absent
- place firewall, filtering, session control functions, reachability constraints, at entry points, rendezvous points
- supports multicast, mobility naturally

Connection State Management

- use virtual links – connections between routers
- static, dynamic virtual links
- trade-off between reachability, state aggregation

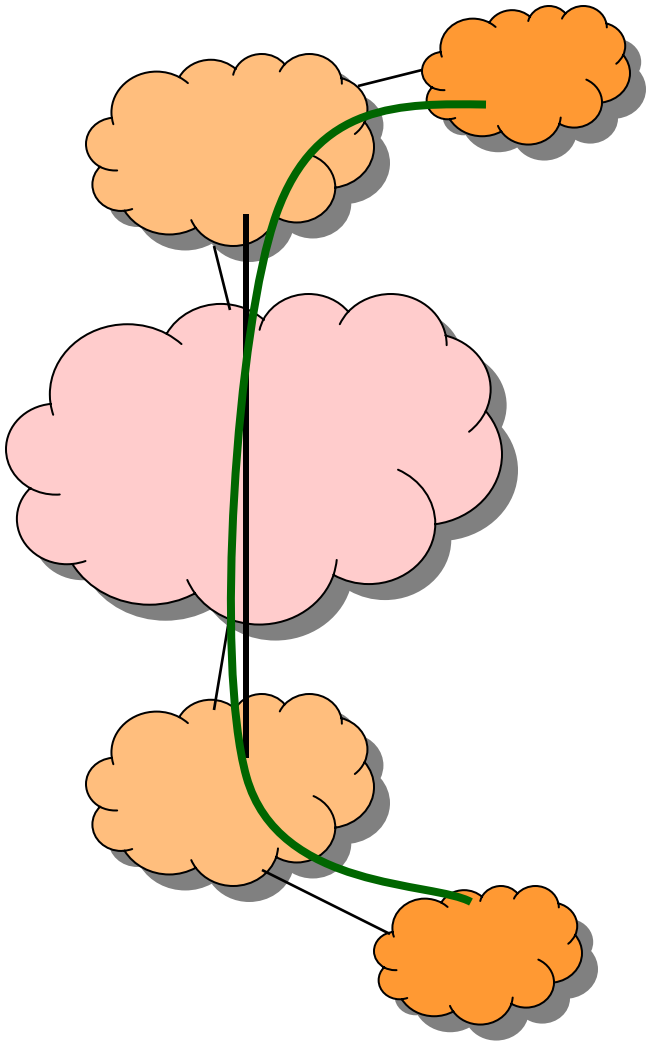

Connection State Management

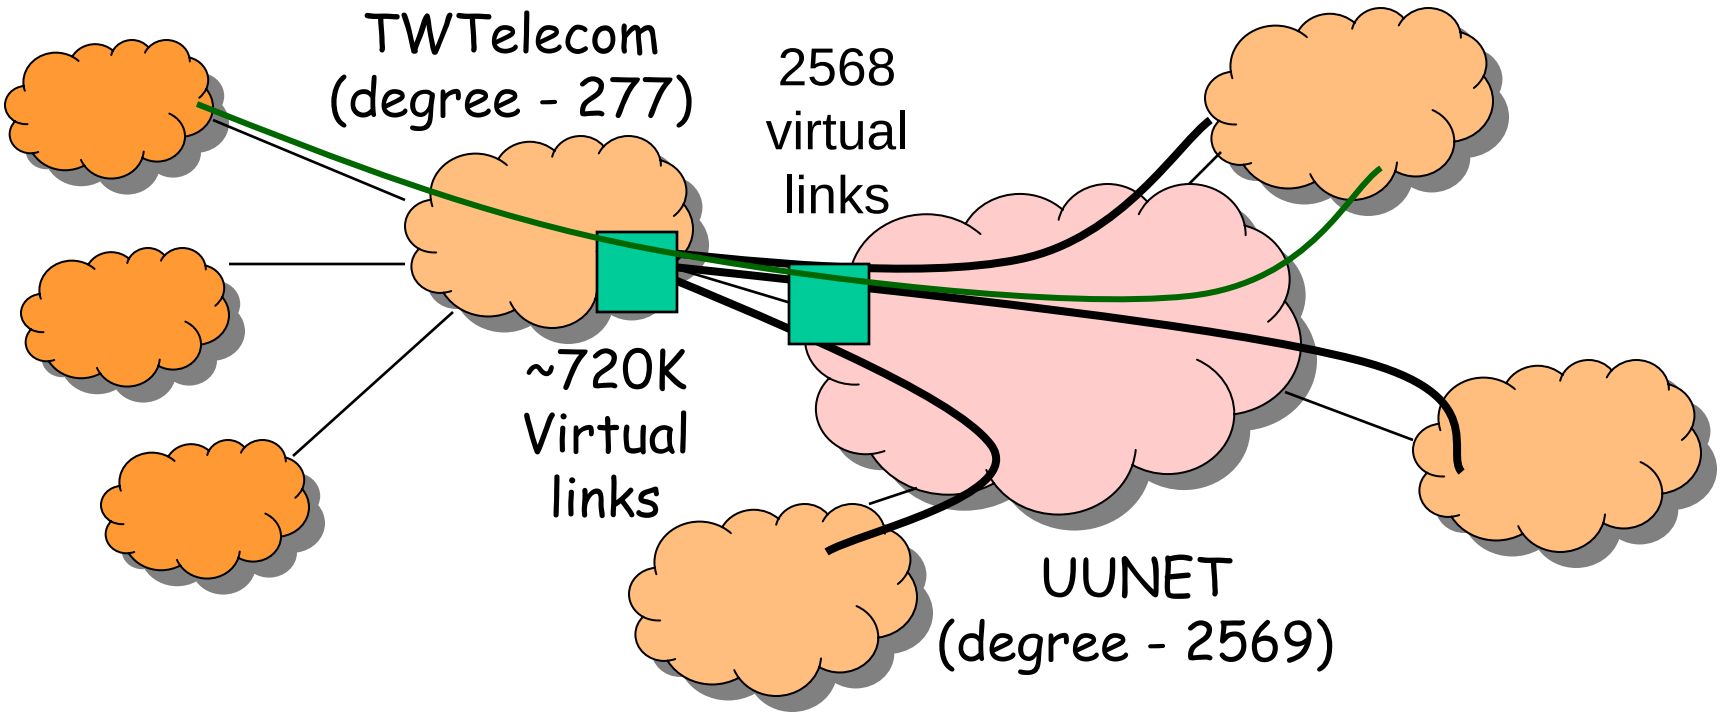

Connection State Management

- place entry points close to clients
- only one connection per entry point => reduced state in middle of network

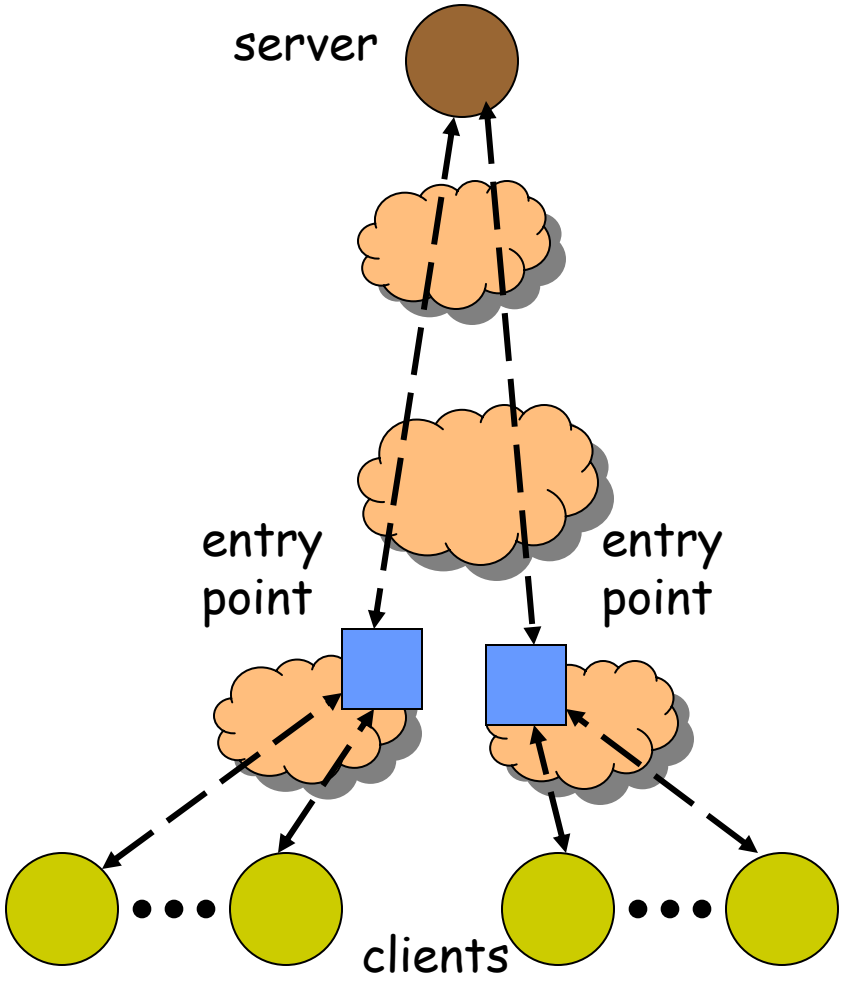

Additional Security Considerations

- reachability only restricted to network layer
- architecture does not prevent user applications from downloading malware
- cannot prevent malware from sending bad traffic to other nodes
- more comprehensive architecture at application layer required

Related Work

- off-by-default (Hotnets 2005)
  - routers avoid keeping routing state for node unless explicitly requested
  - any change in node's decision must be propagated throughout network
  - large number of messages, large delays
- i3 (Sigcomm 2002)
  - indirection very similar to our use of entry, rendezvous points
  - overlay solution on top of IP
  - does not address client reachability

Conclusions

- new connection oriented architecture
- high level ideas only, still work-in-progress
- (we believe) architecture is viable, necessary for aiding Internet security

More Related Work (Backup)

- SOS (Sigcomm'02)
  - critical servers can only be reached through certain special nodes
  - but servers have well-known IP addresses that can be leaked out, all routers around servers must drop packets not from special nodes
  - does not address client reachability
- DoS-resistance (FDNA'04)
  - separate private client address space
  - return path (static domain id) appended to packets towards servers
  - once domain ids known, attack packets can be sent to clients