The Negotiation Problem

Bryan Ford

Yale University

bryan.ford@yale.edu

Jana Iyengar

Franklin & Marshall College

jiyengar@fandm.edu

Presentation for TSVAREA meeting – Nov 13, 2009

http://baford.info/tng

A Proliferation of Layers and Layer Combinations

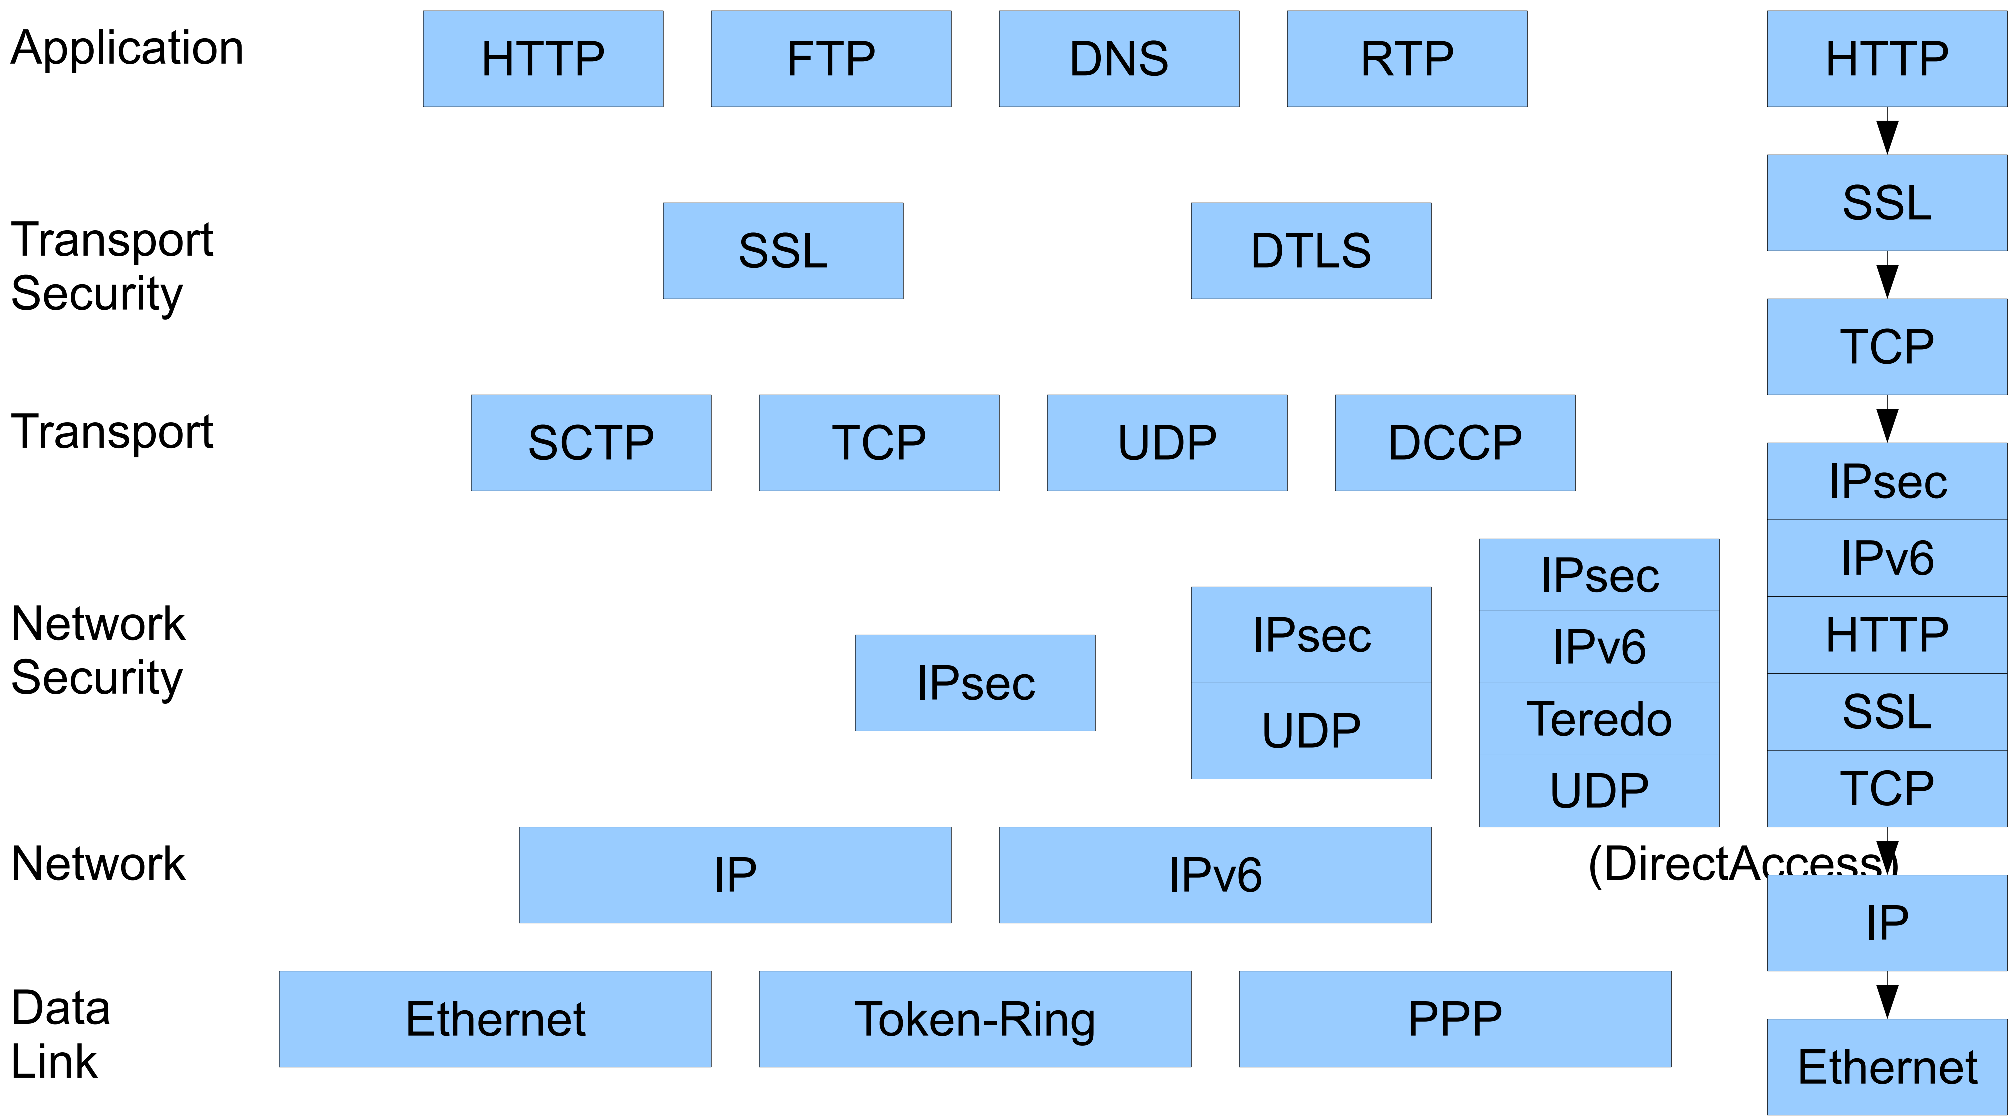

Future: Ever More Layers/Combinations?

Multi-Streaming Transports
SCTP [rfc4960],
SST [SIGCOMM'07]

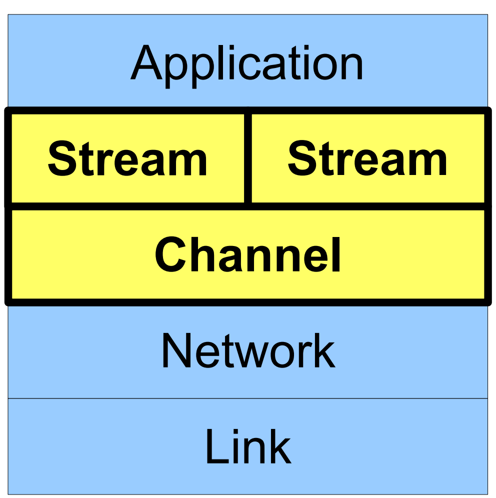

Multipath Transports
SCTP [rfc4960],
MPTCP [WIP]

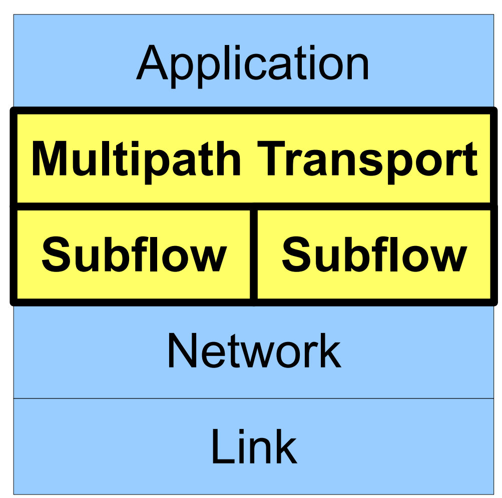

Further Decomposition
[“Breaking Up the Transport Logjam”, HotNets'08]

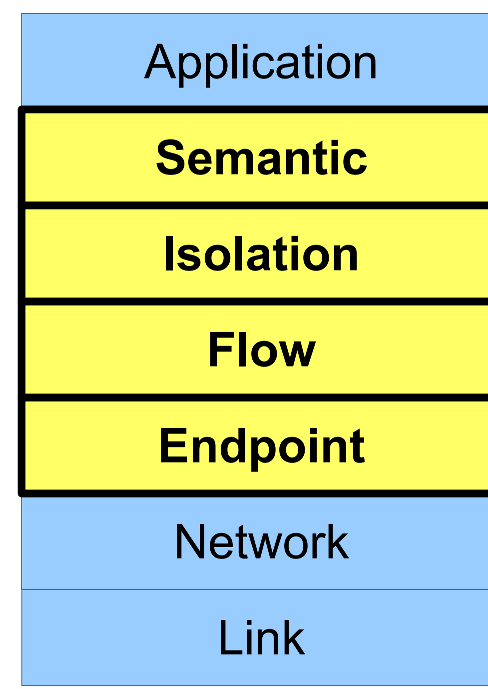

The Negotiation Problem

Decisions, decisions!

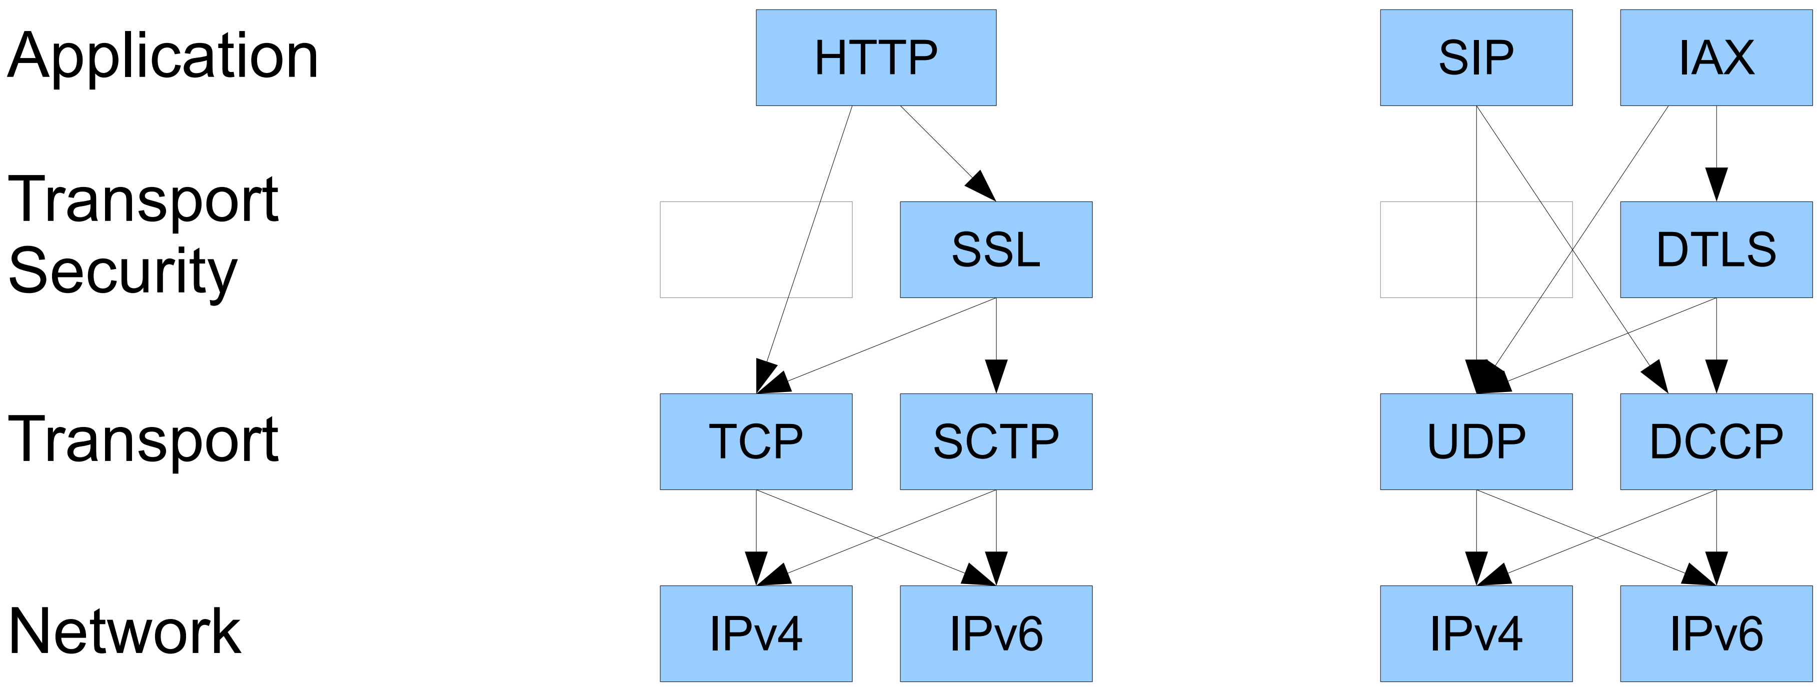

Compatibility and Preference

Which combinations do *both* endpoints support?

Which combinations do they *prefer*?

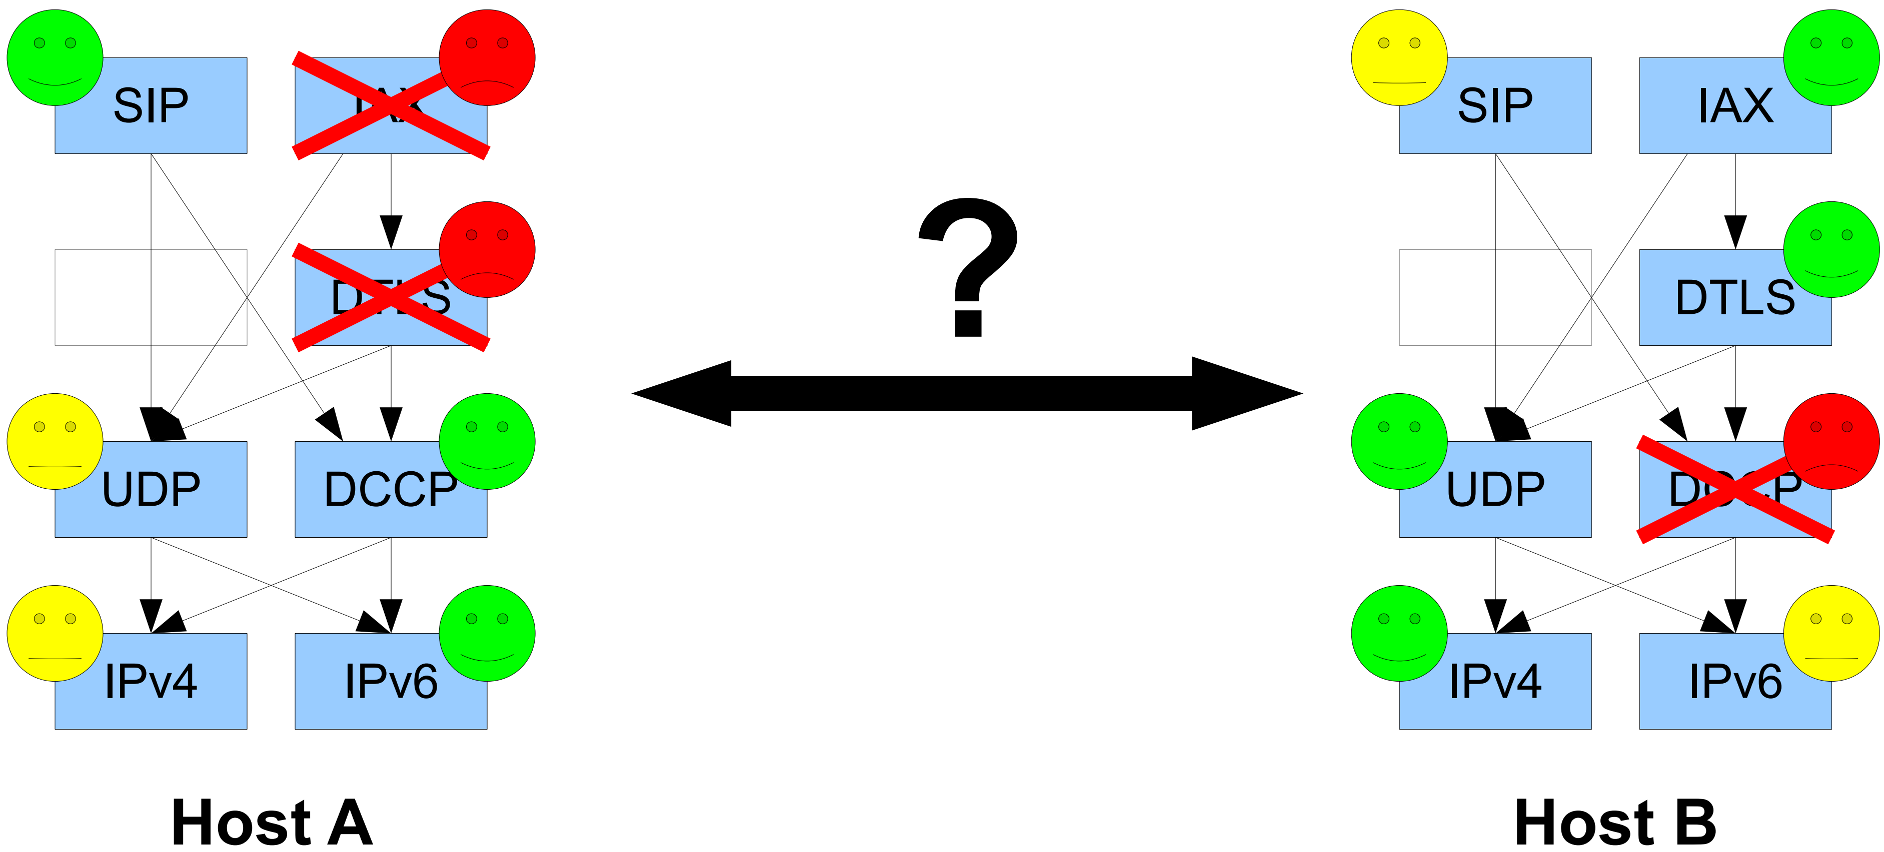

Talk Outline

- Three negotiation strategies (2 explicit, 1 implicit)
  - Including a new in-band negotiation mechanism
  - Combined explicit/implicit negotiation
- A framework for negotiation
- Discussion

Negotiation Strategies

Approach 1: Try and Fall Back

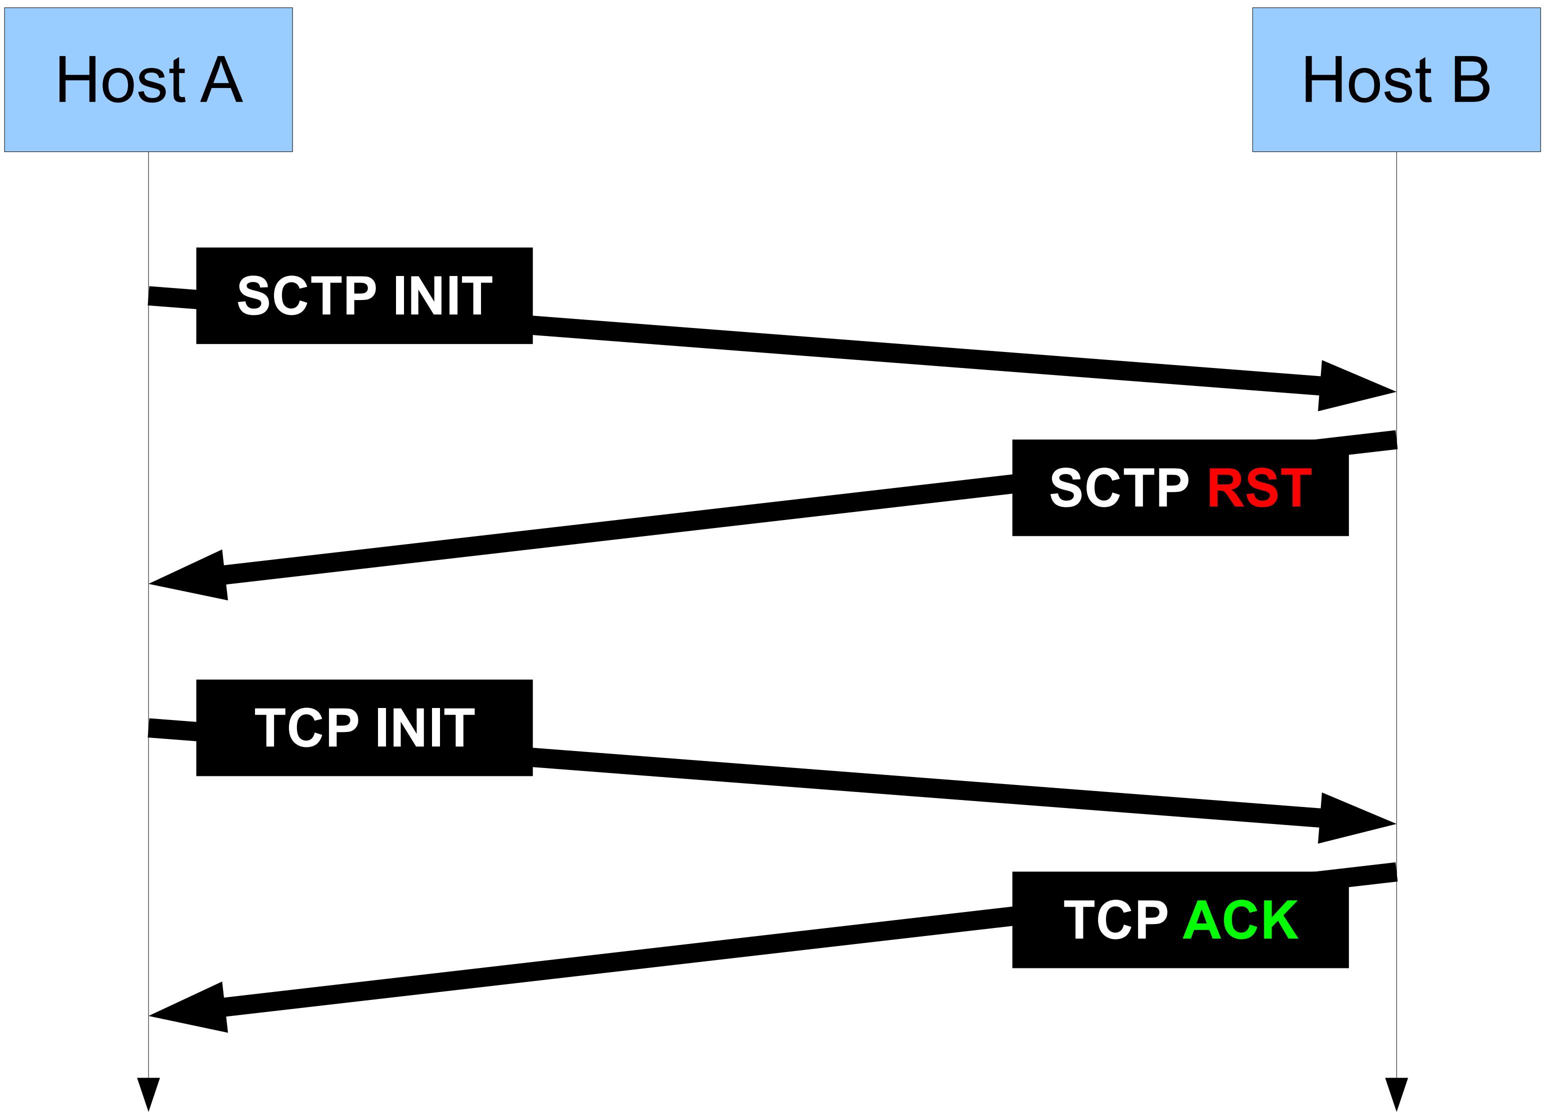

Challenge 1: Controlling Delay

- Failures can incur *timeouts* (e.g., due to NATs)
- … potentially *compounded* by layering

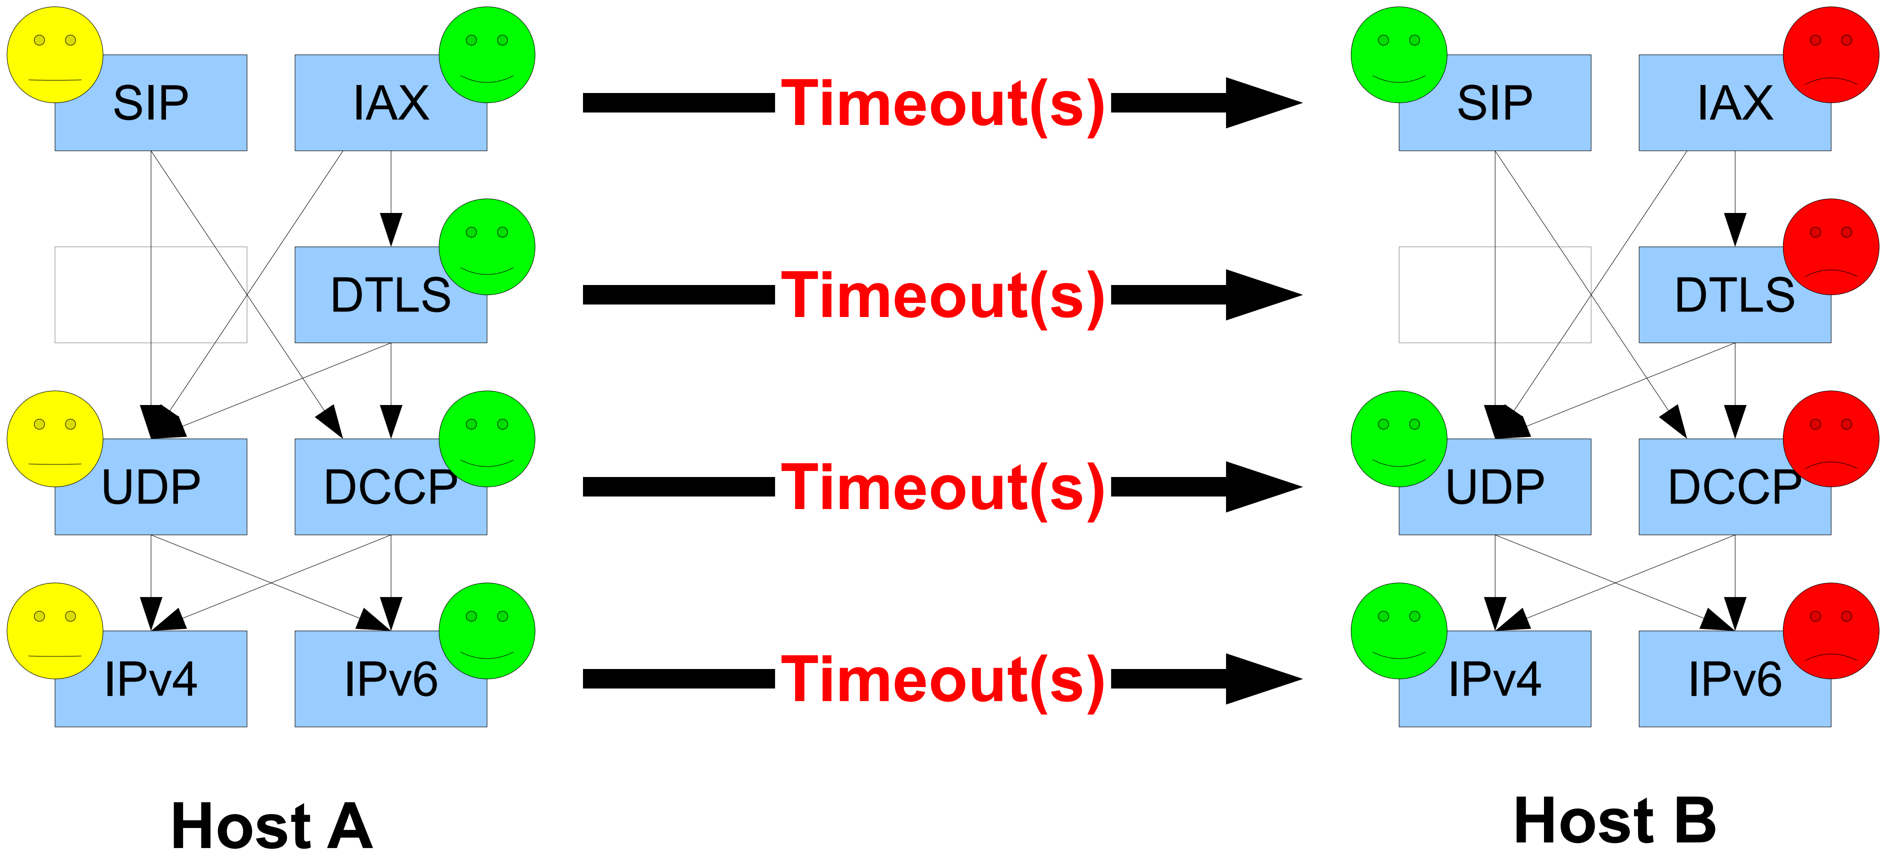

Approach 2: Try in Parallel

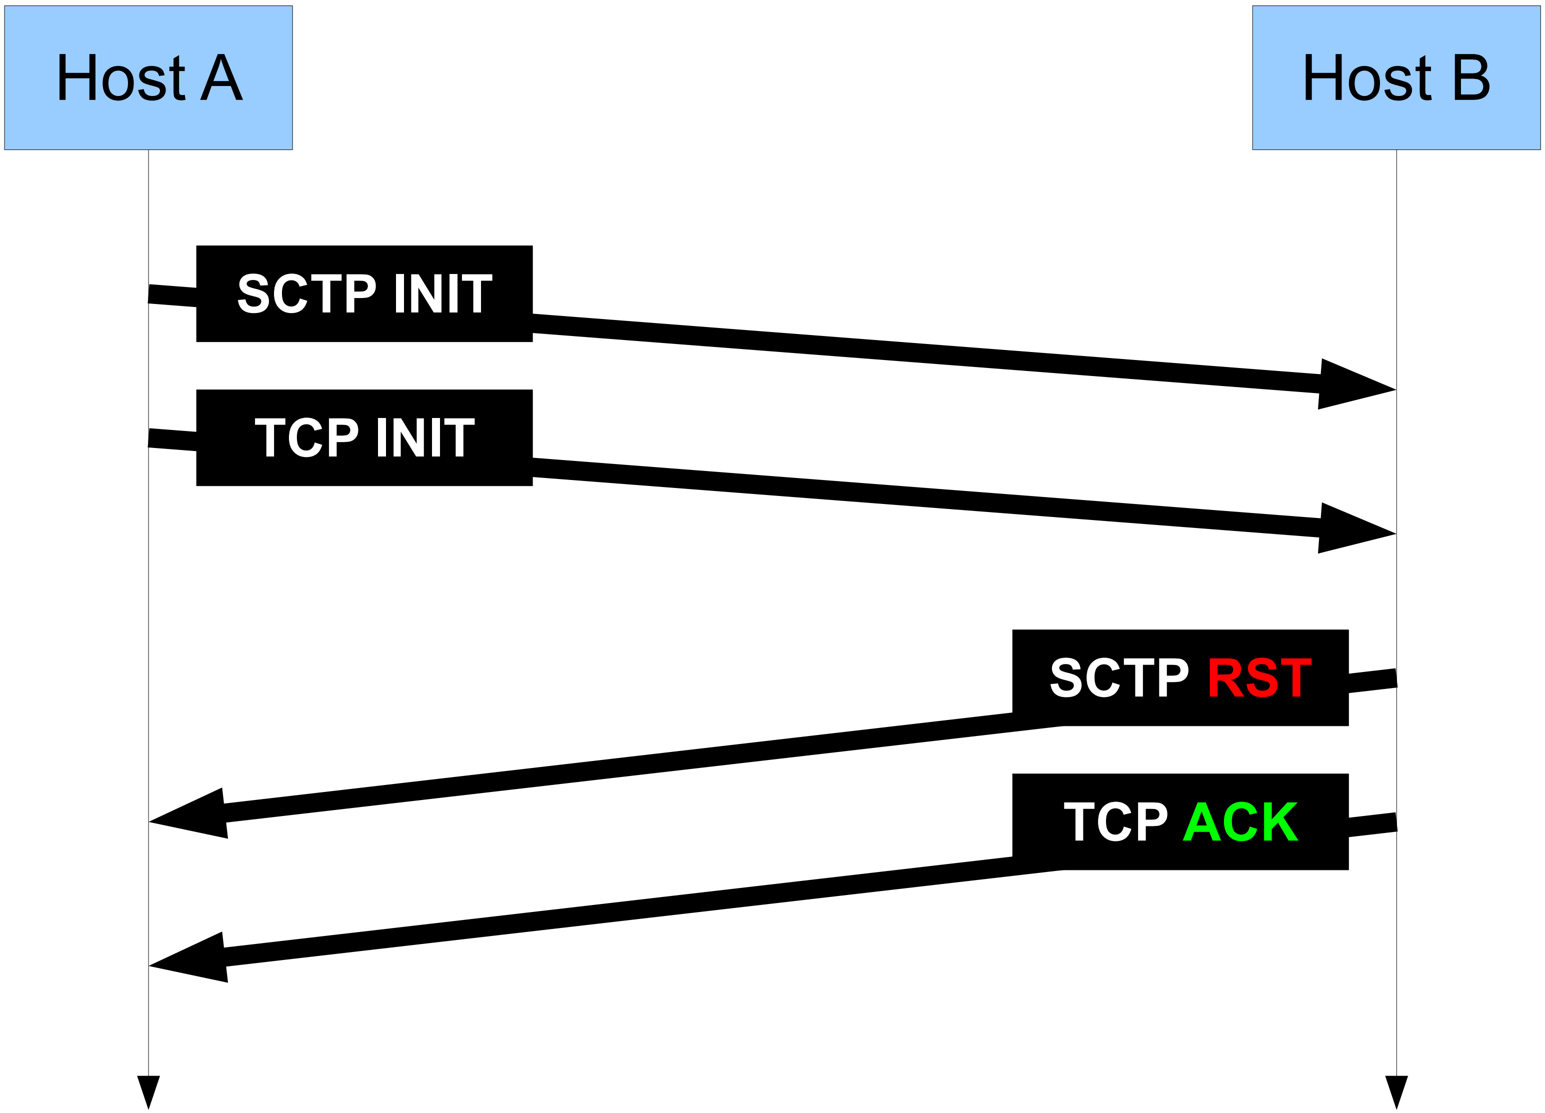

Challenge 2a: Redundant State

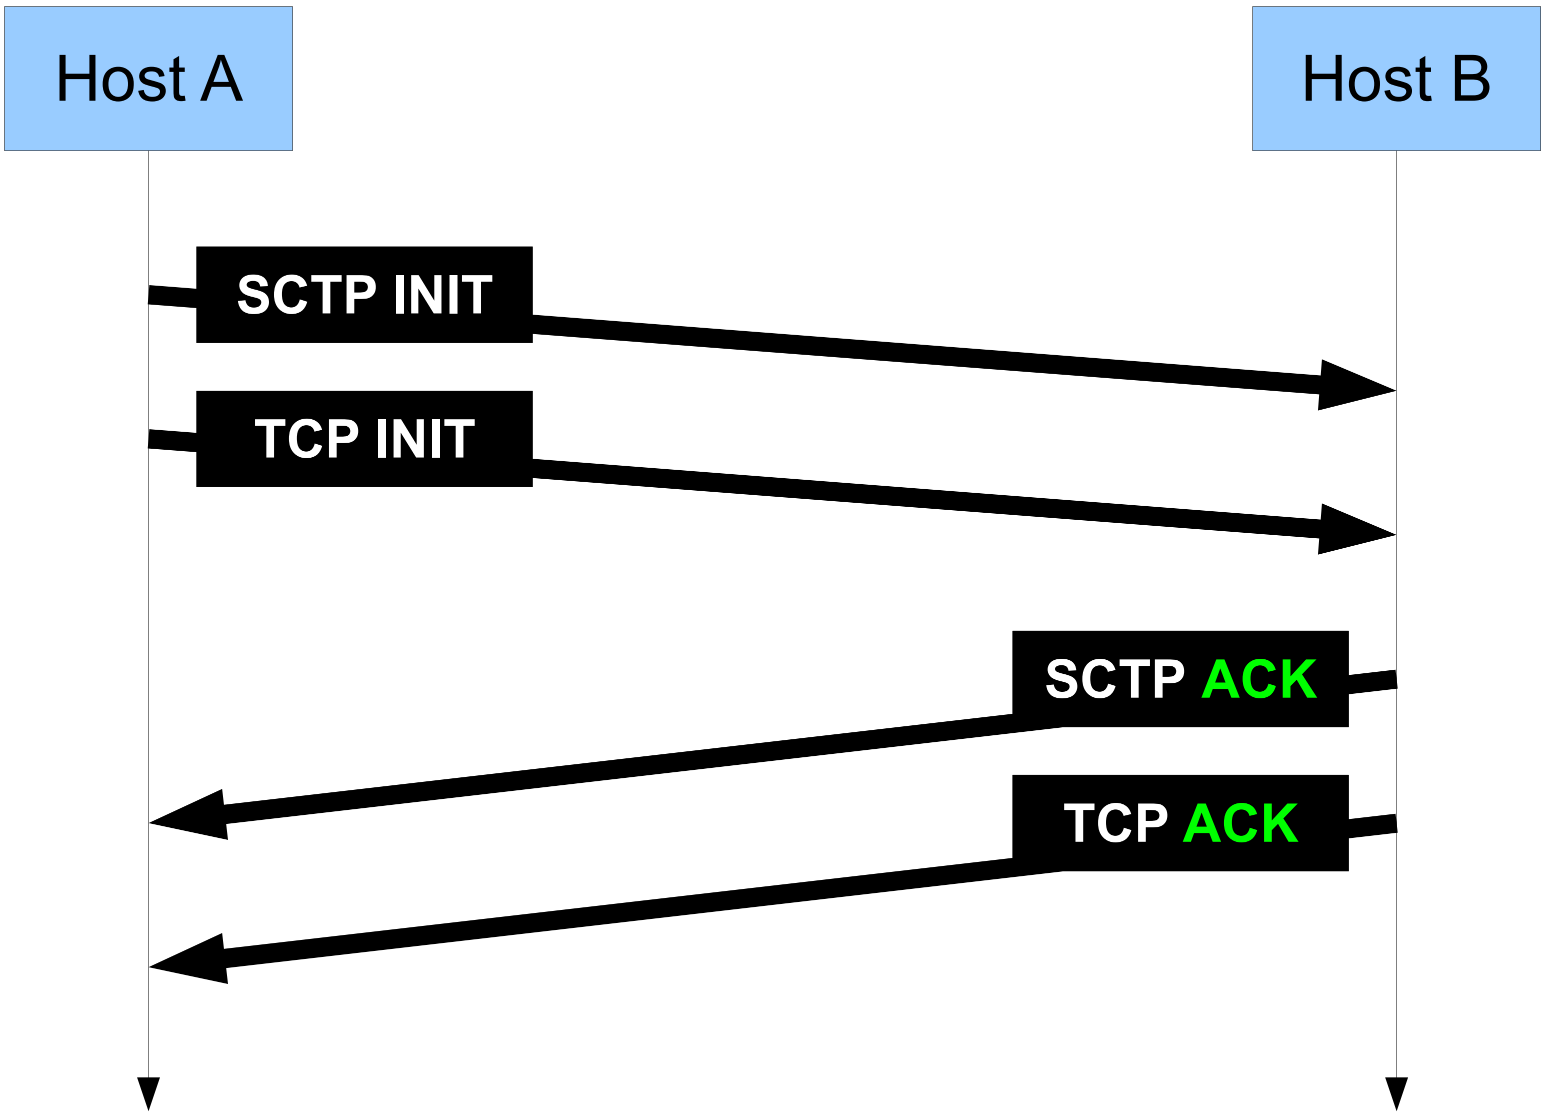

Challenge 2b: Combinations

Layering can lead to explosion of choices

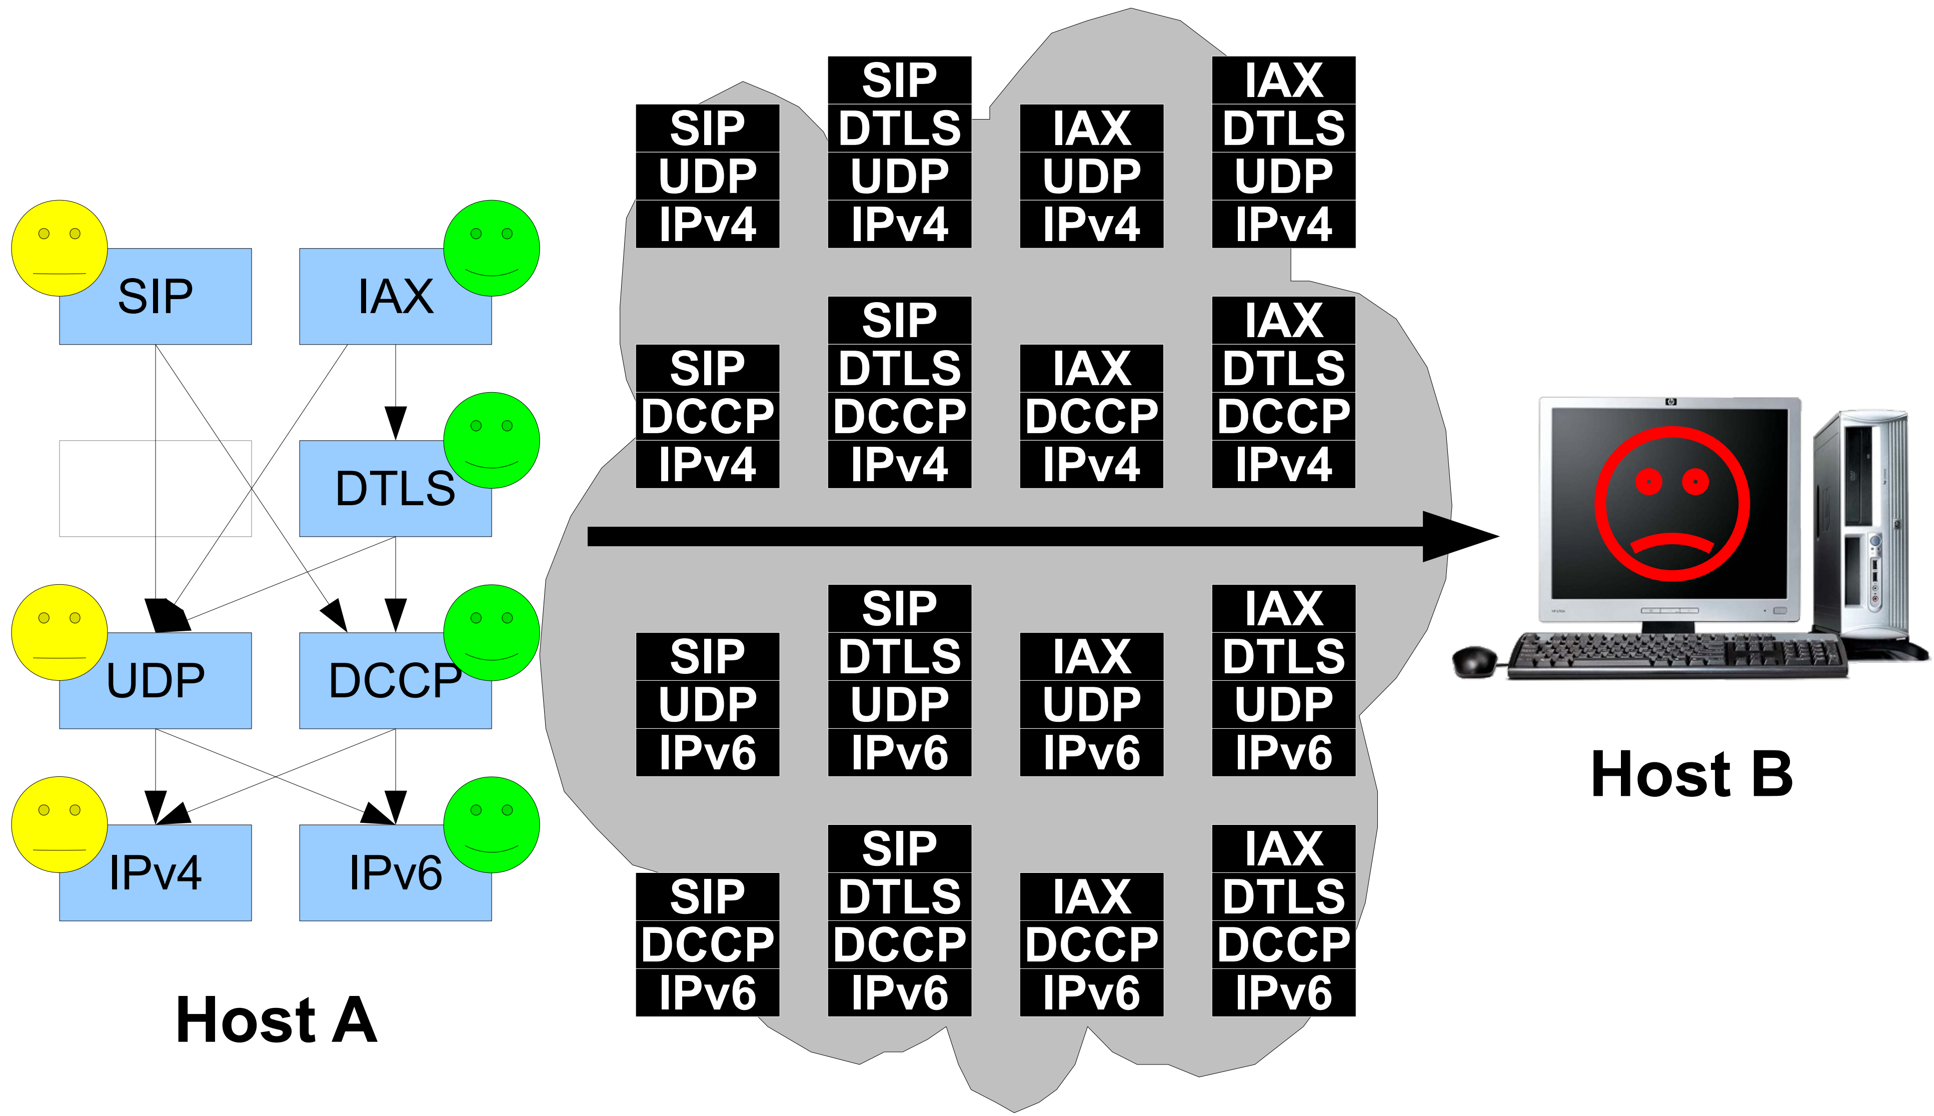

Negotiation Strategies

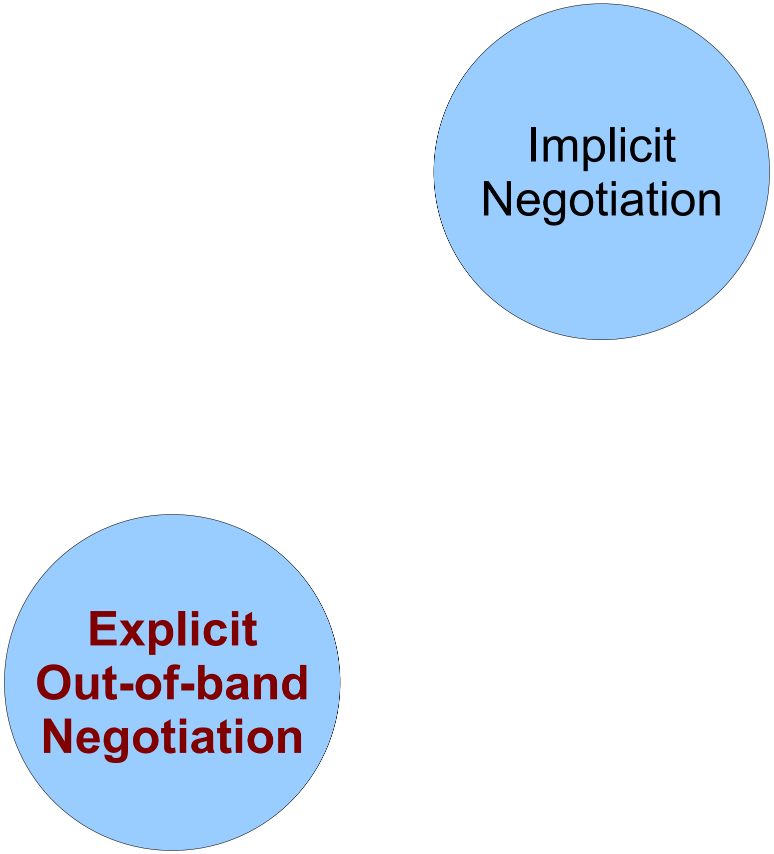

Approach 3: Out-of-Band Information

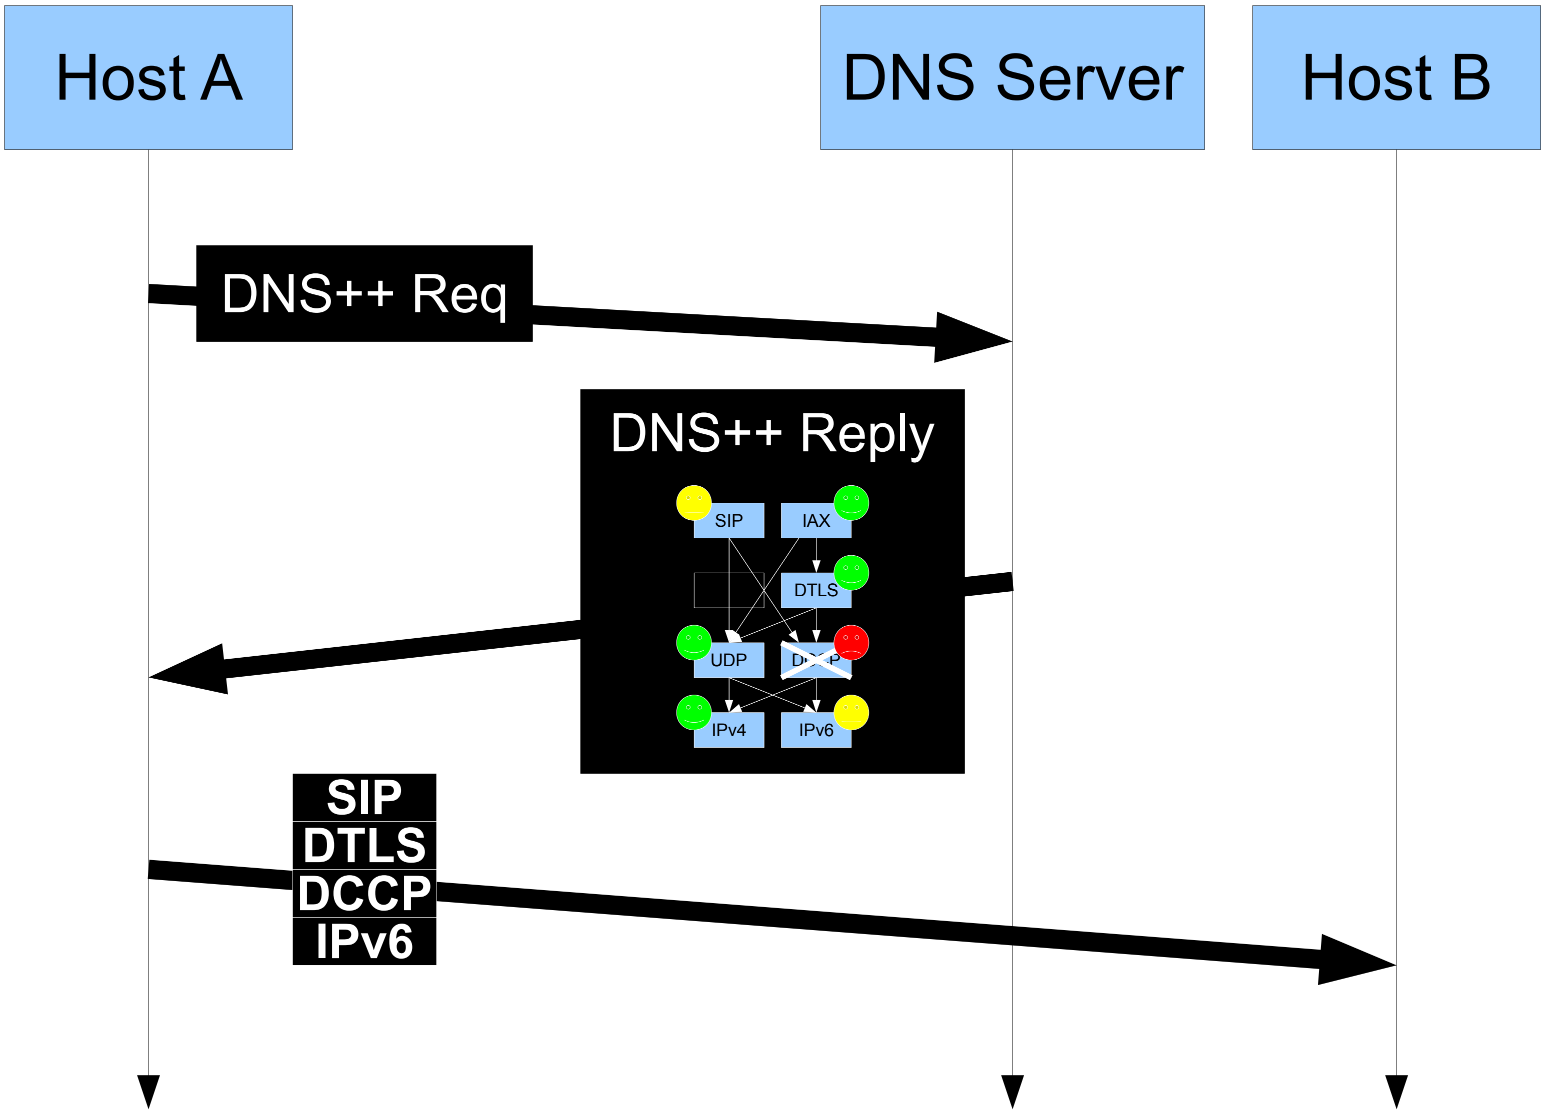

Challenge 3a: Administration

DNS server must know:

- Name → IP mapping (as before)
- Entire protocol stack supported by Host B
- Protocol options?

⟹ Synchronization Nightmare?

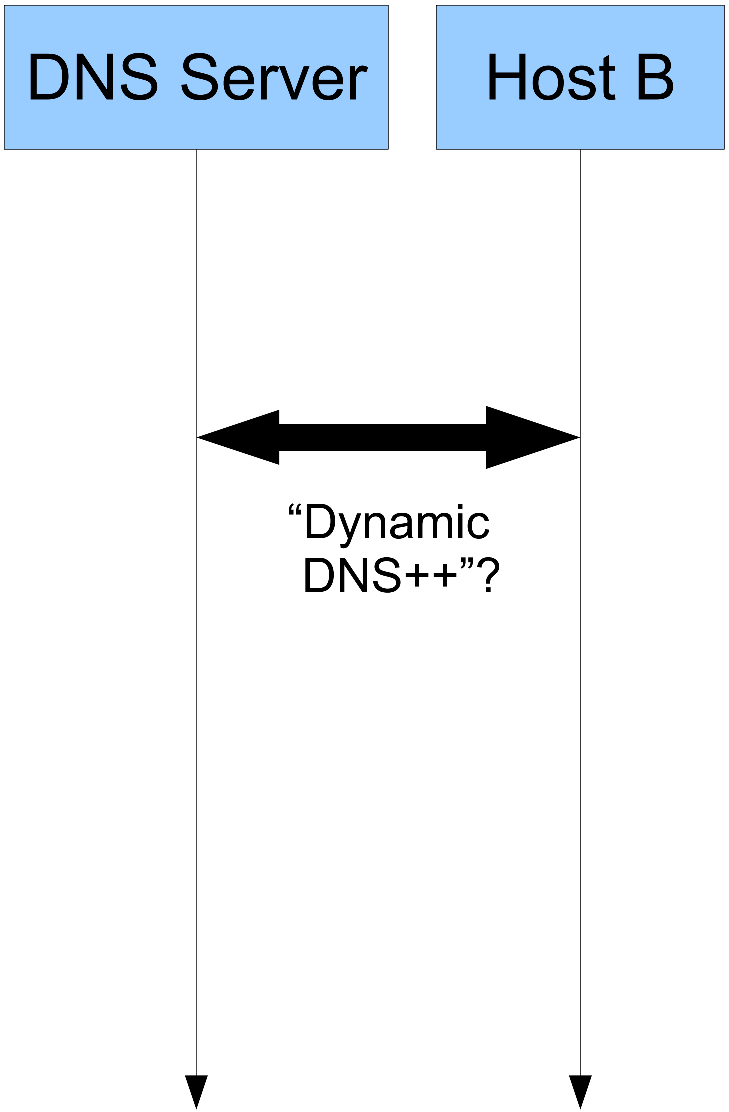

Challenge 3b: E2E Robustness

If endpoints agree on configuration X, *will it work?*

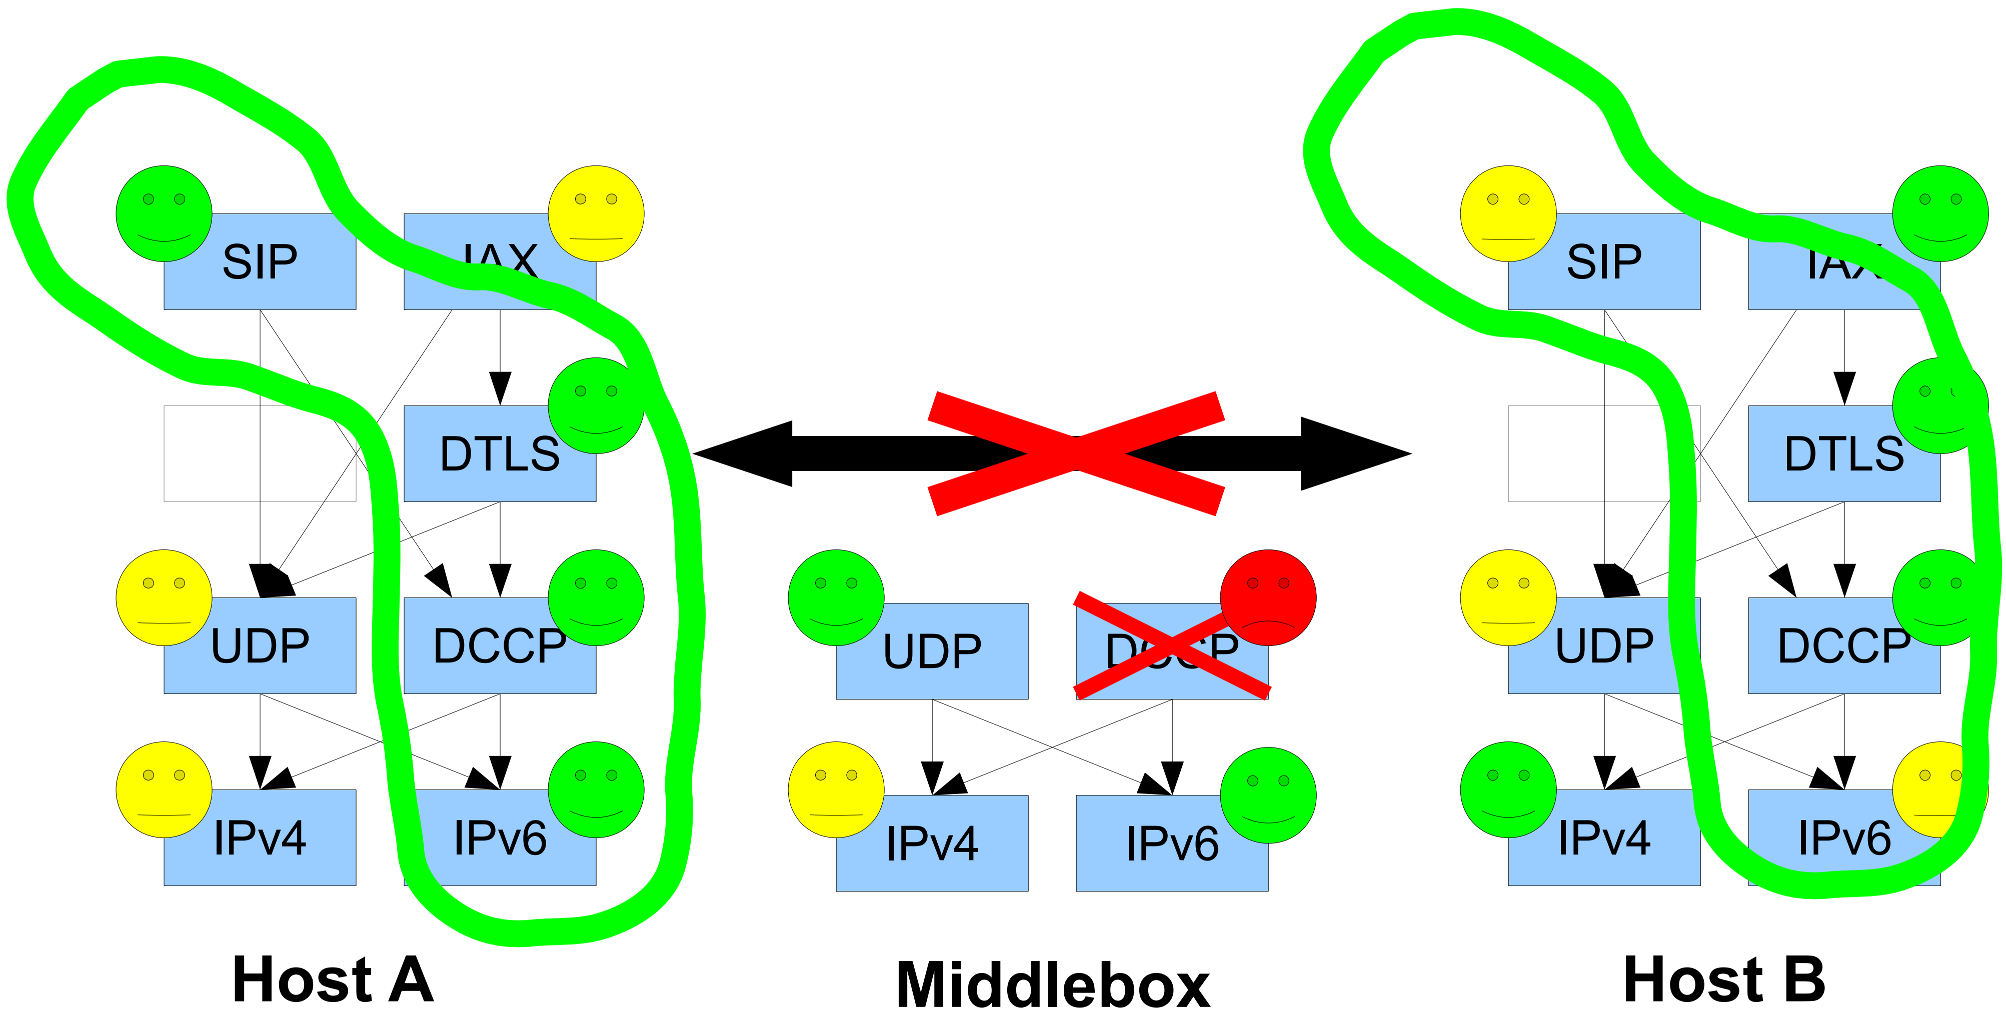

Negotiation Strategies

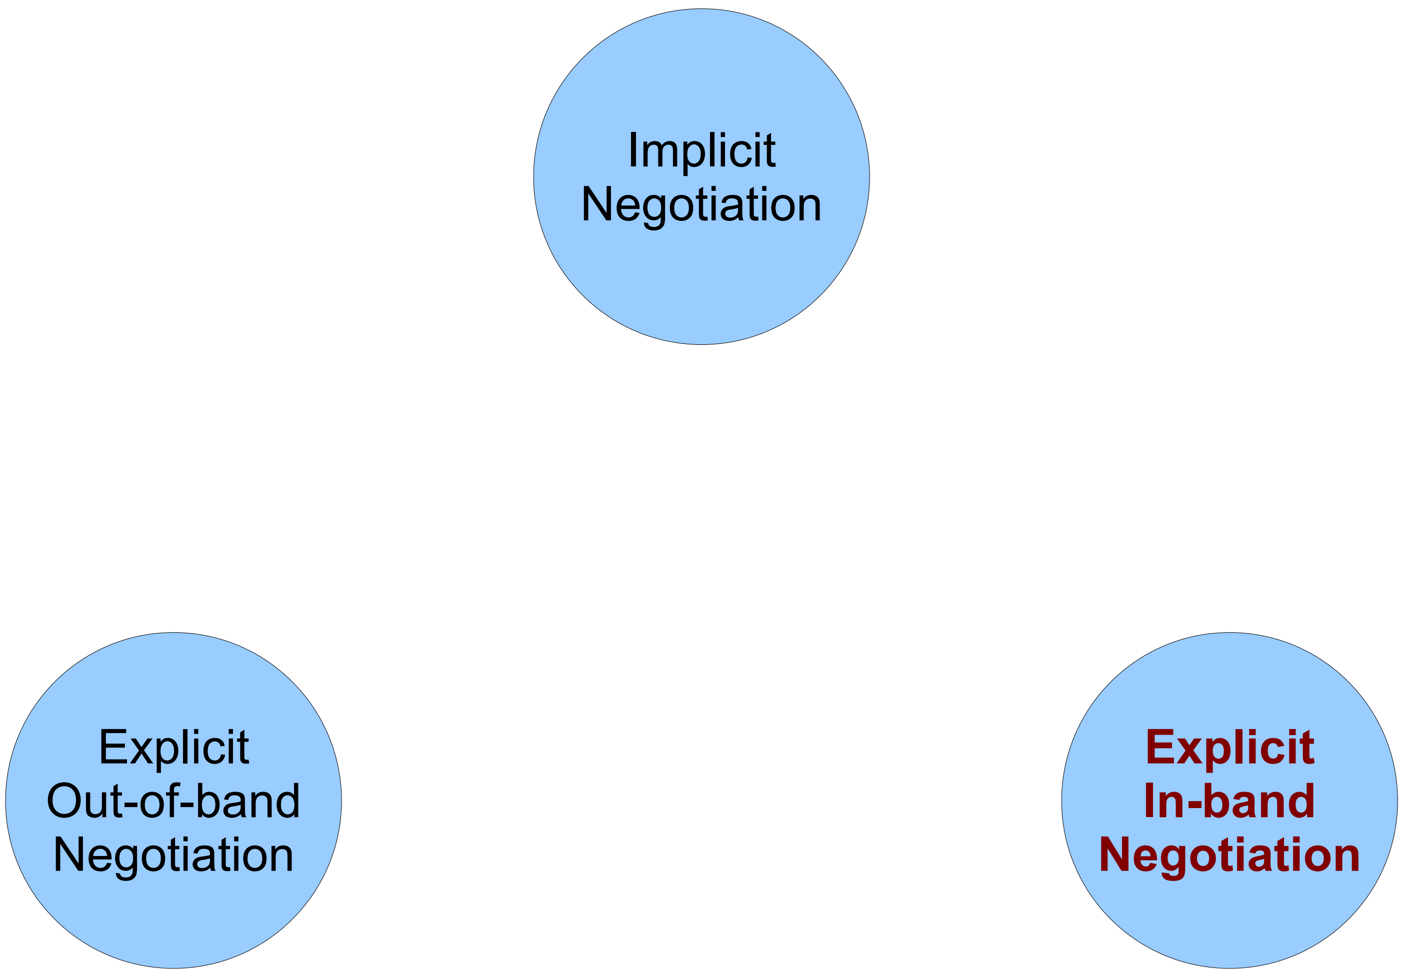

Approach 4: In-band Negotiation

- Hosts explicitly describe possible configurations during initial "meta-communication" exchange, *before* actual communication commences

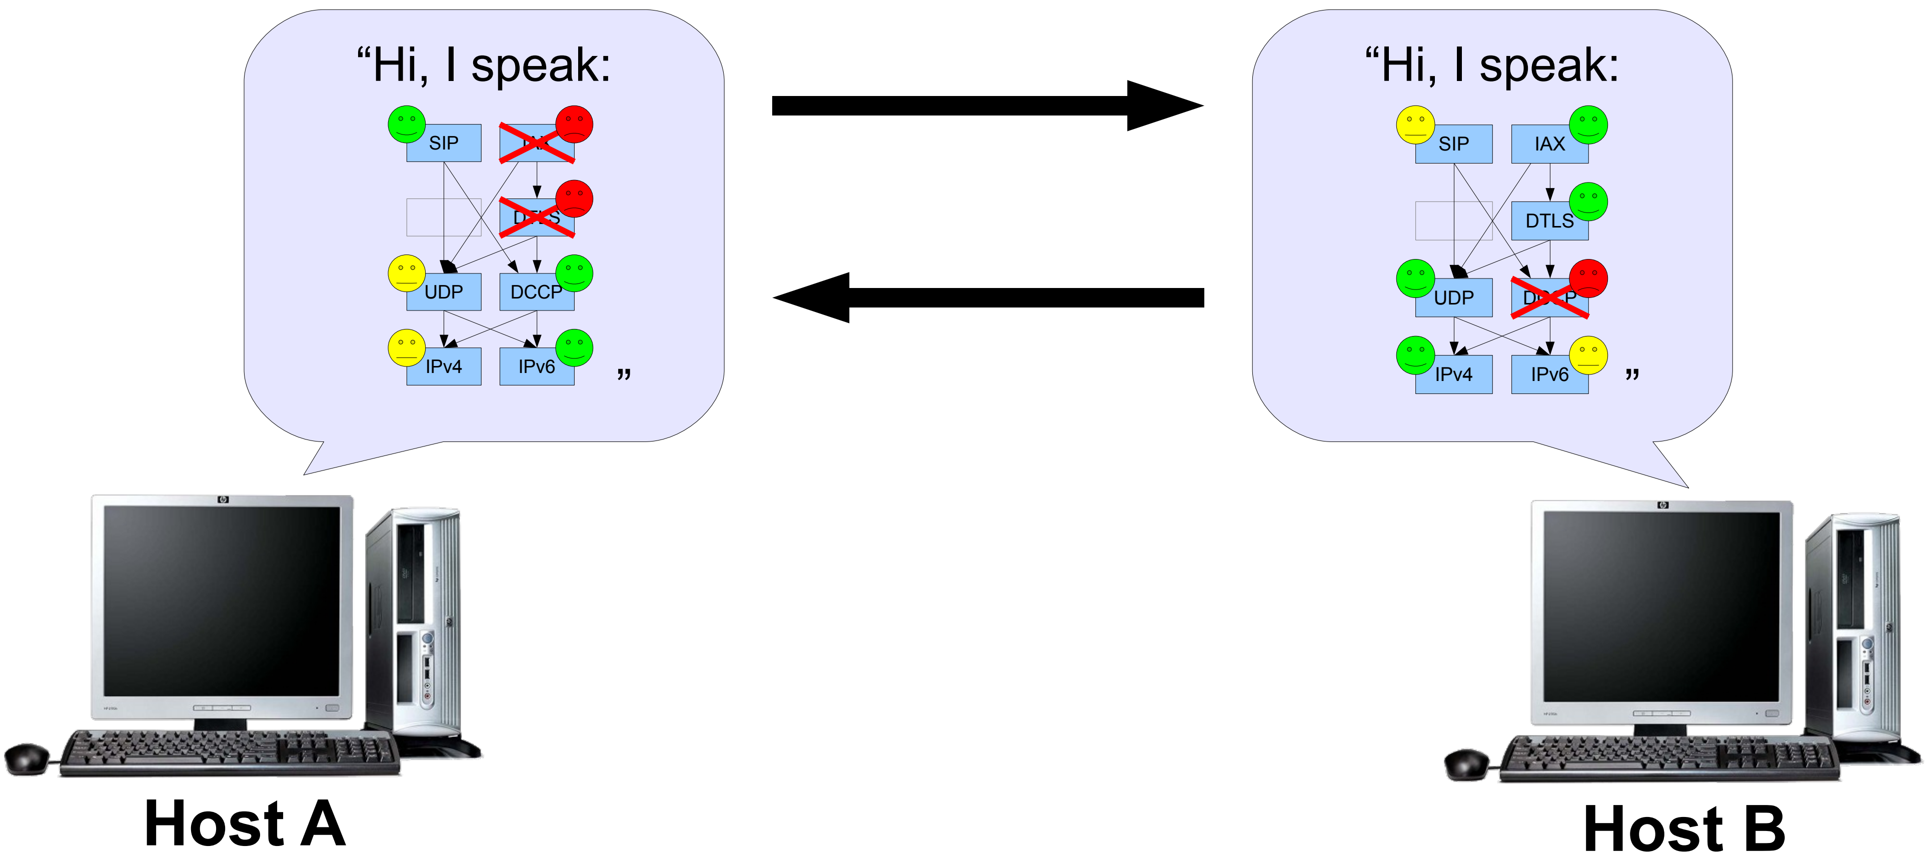

Message 1: Initiator → Responder: **Propose Protocol Graph**

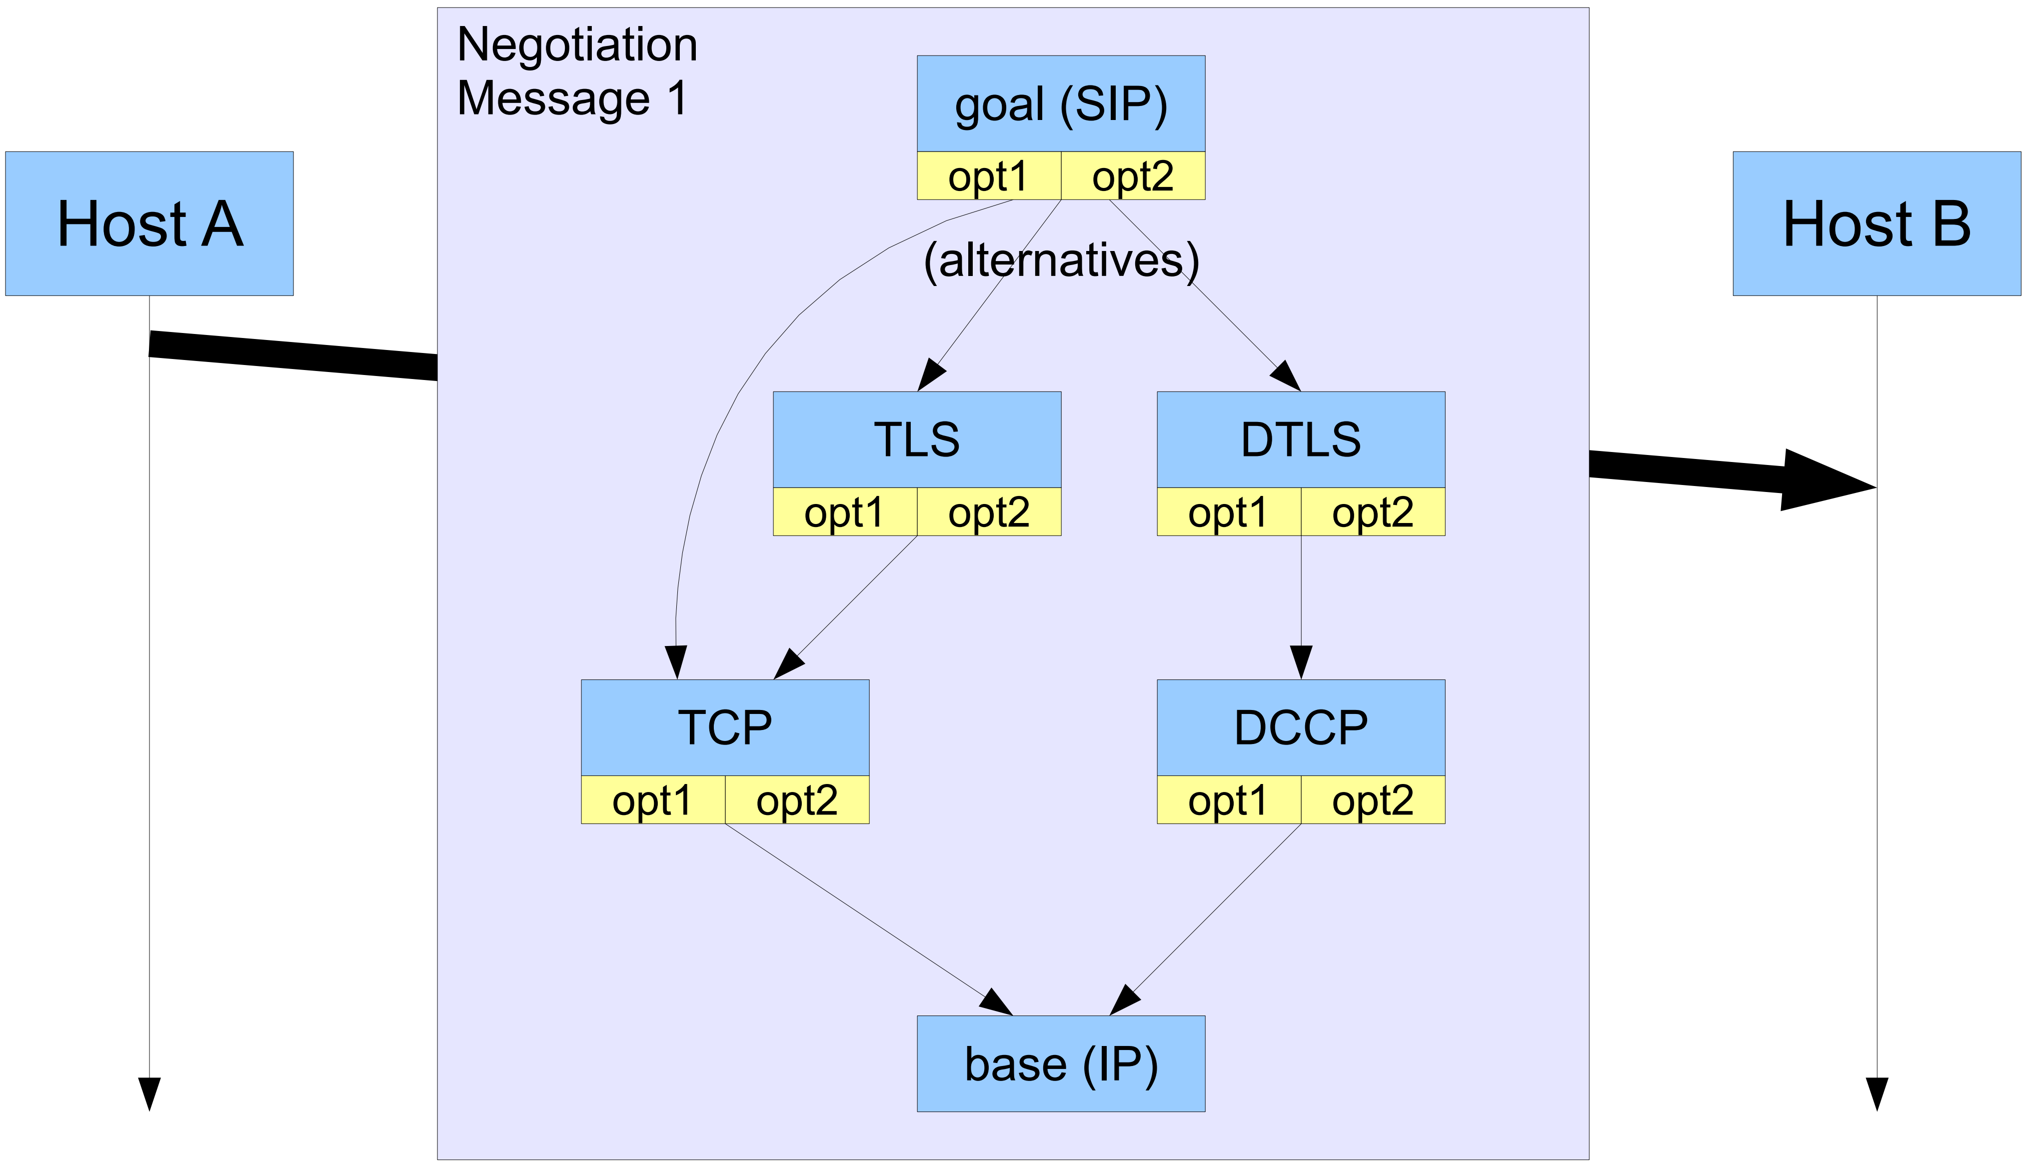

Message 2: Responder → Initiator: **Revise Protocol Graph**

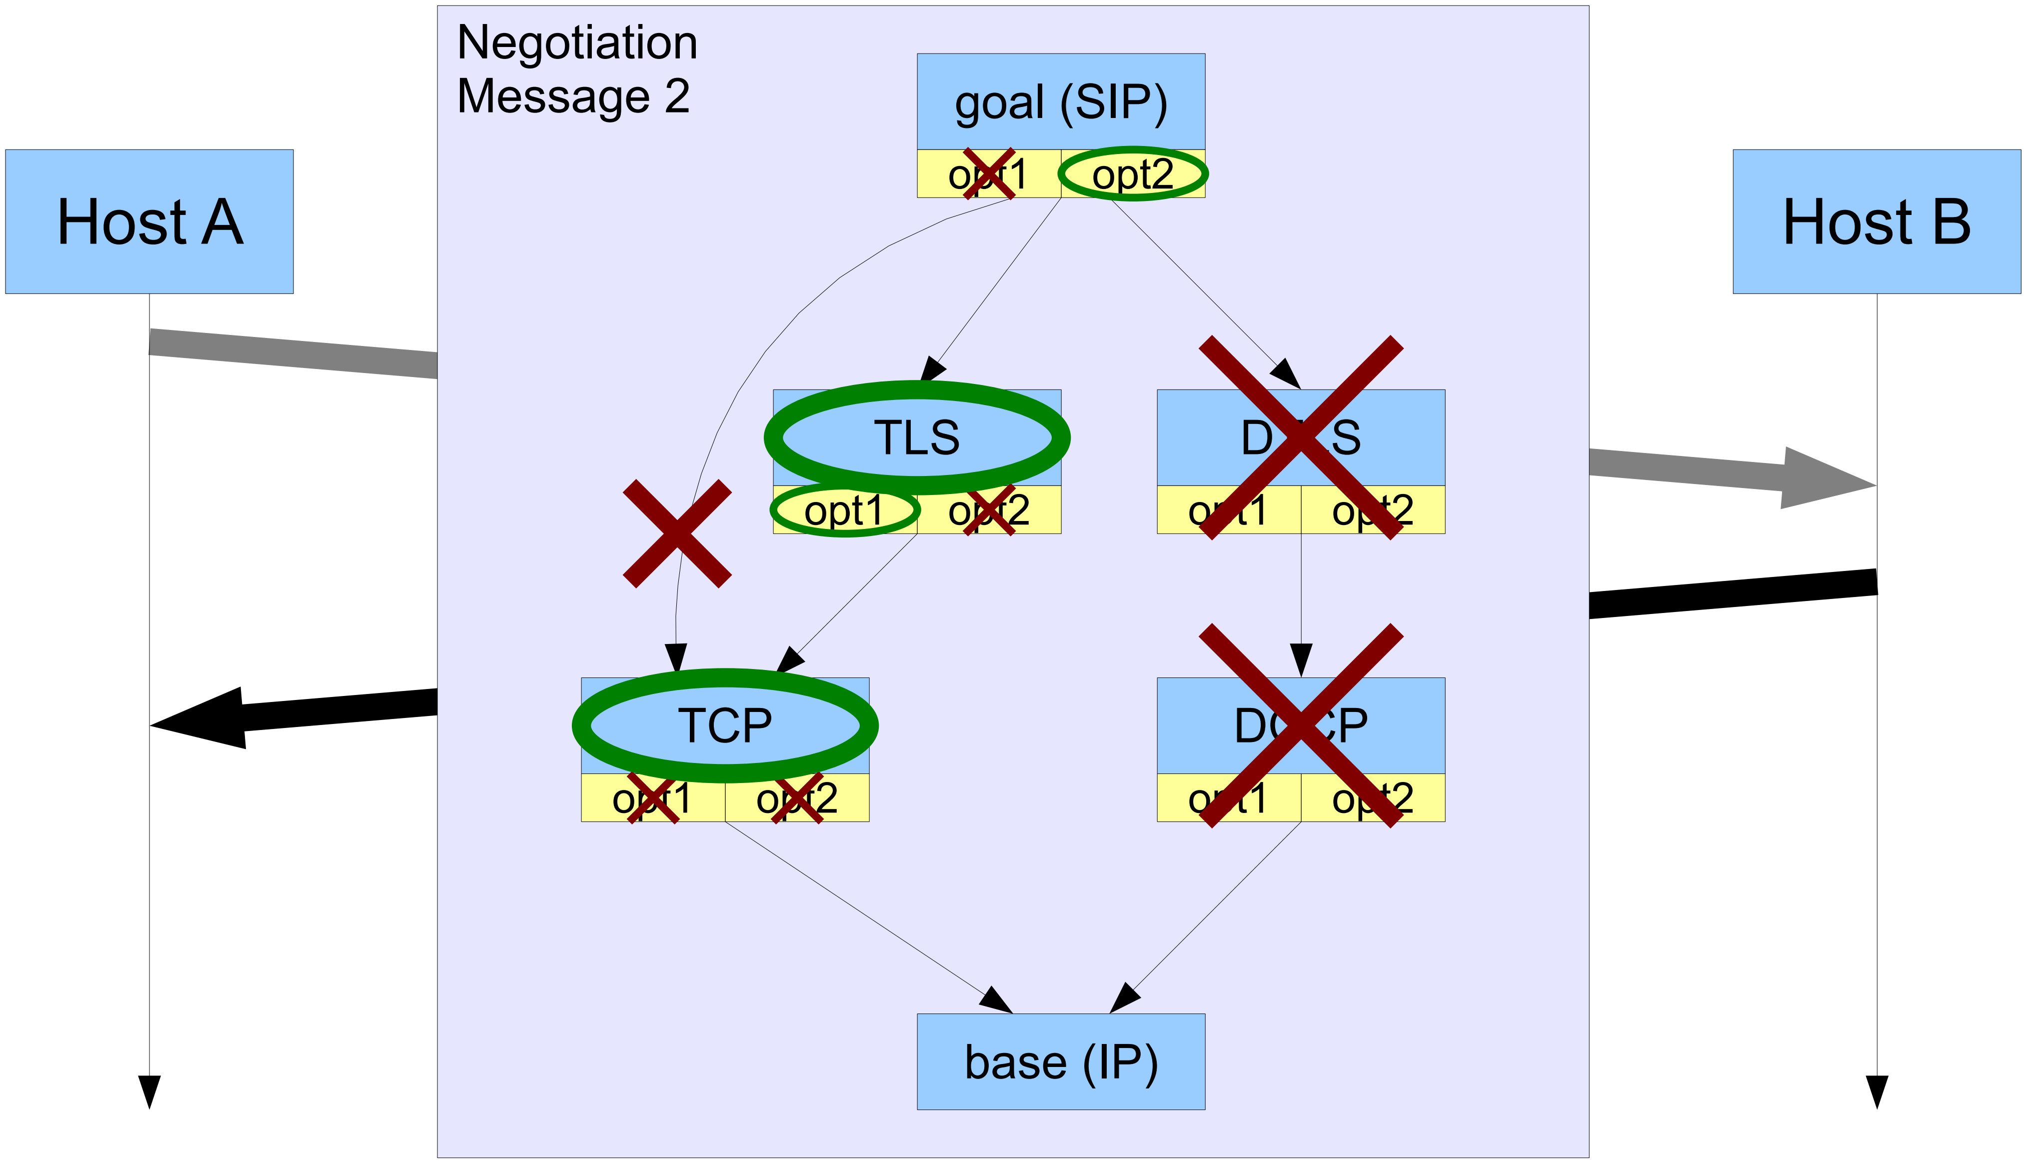

Message 3: Initiator → Responder: **Acknowledge Protocol Graph**

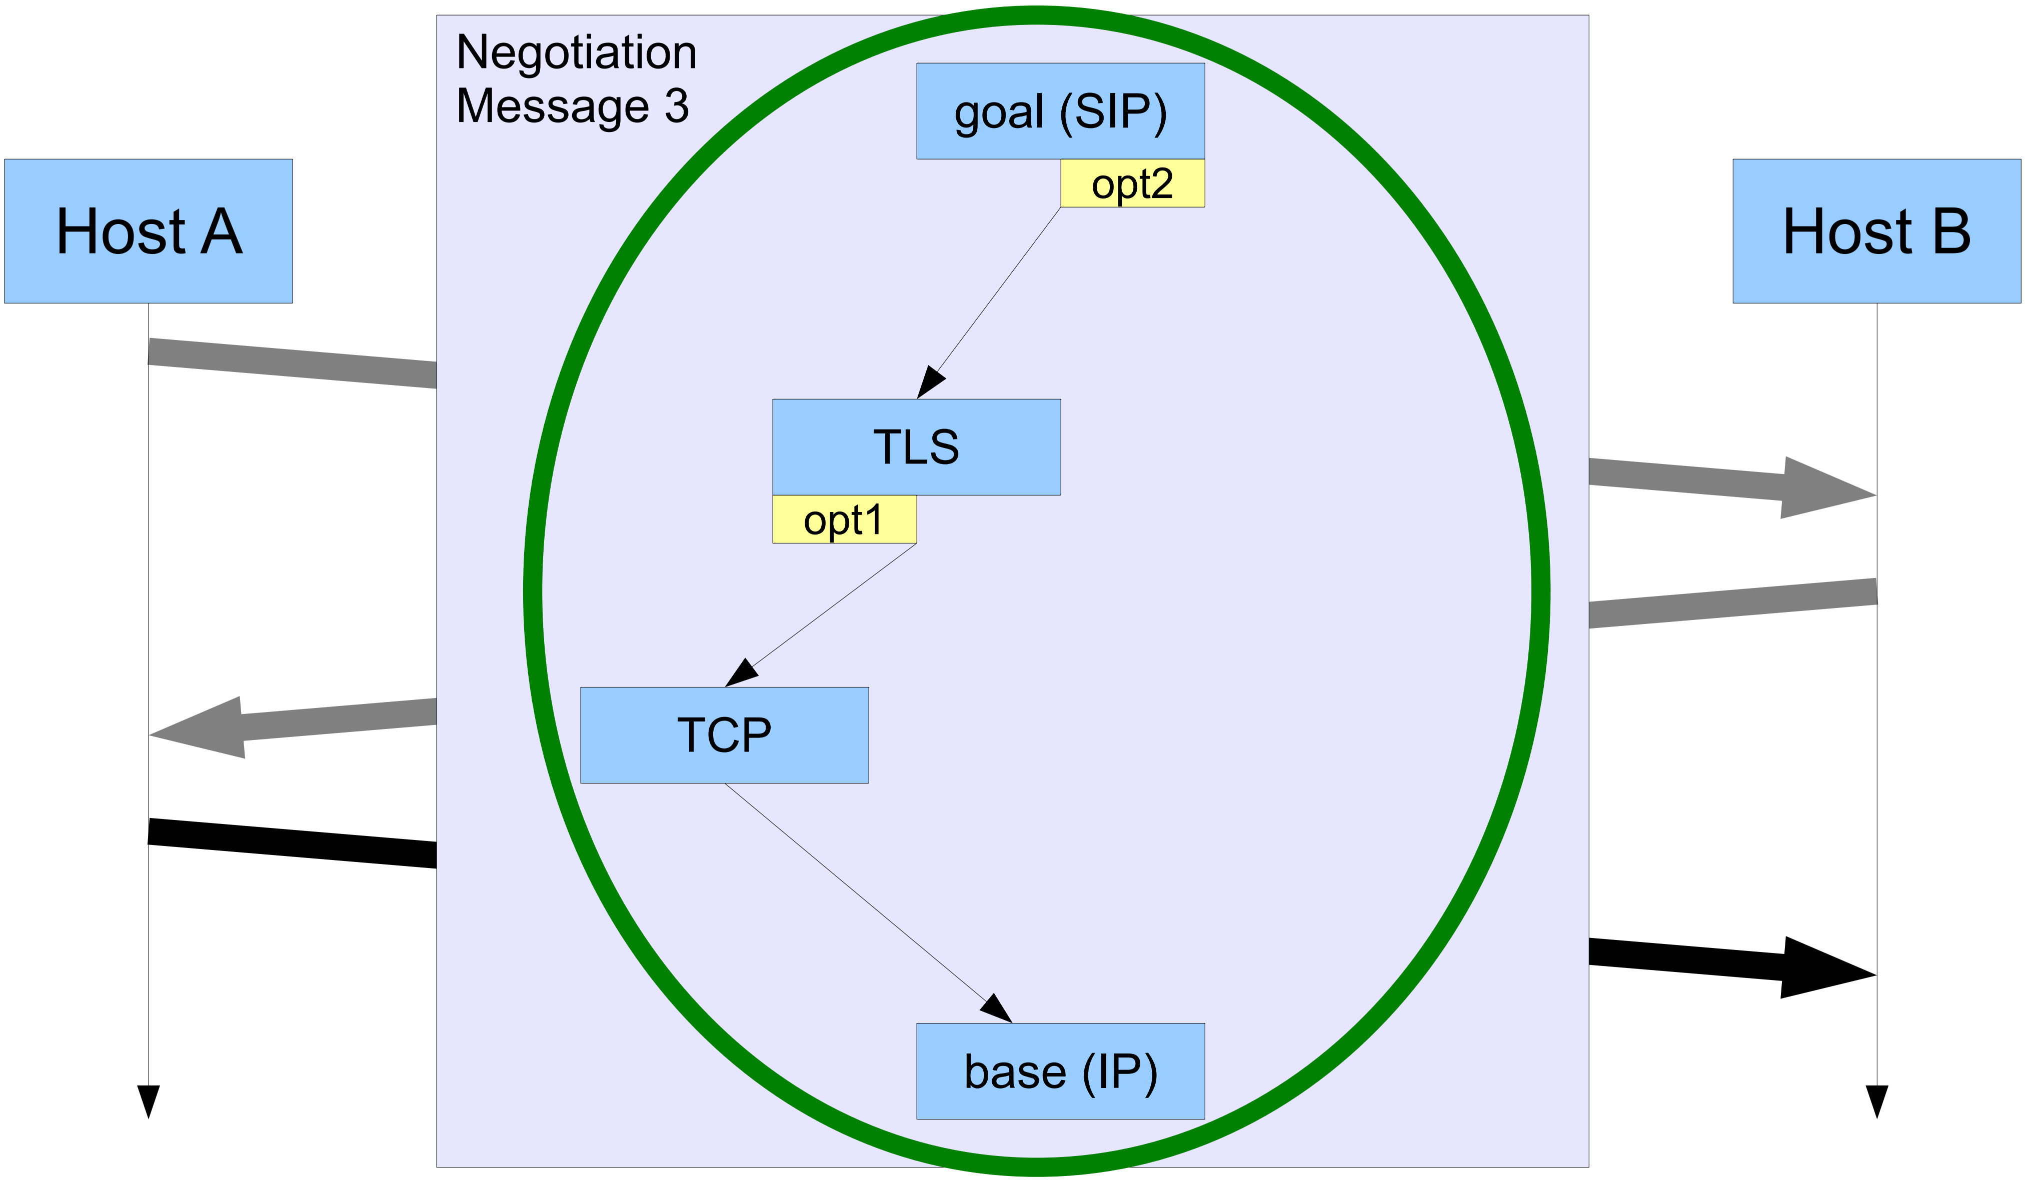

Message 4+:
According to Negotiated Stack

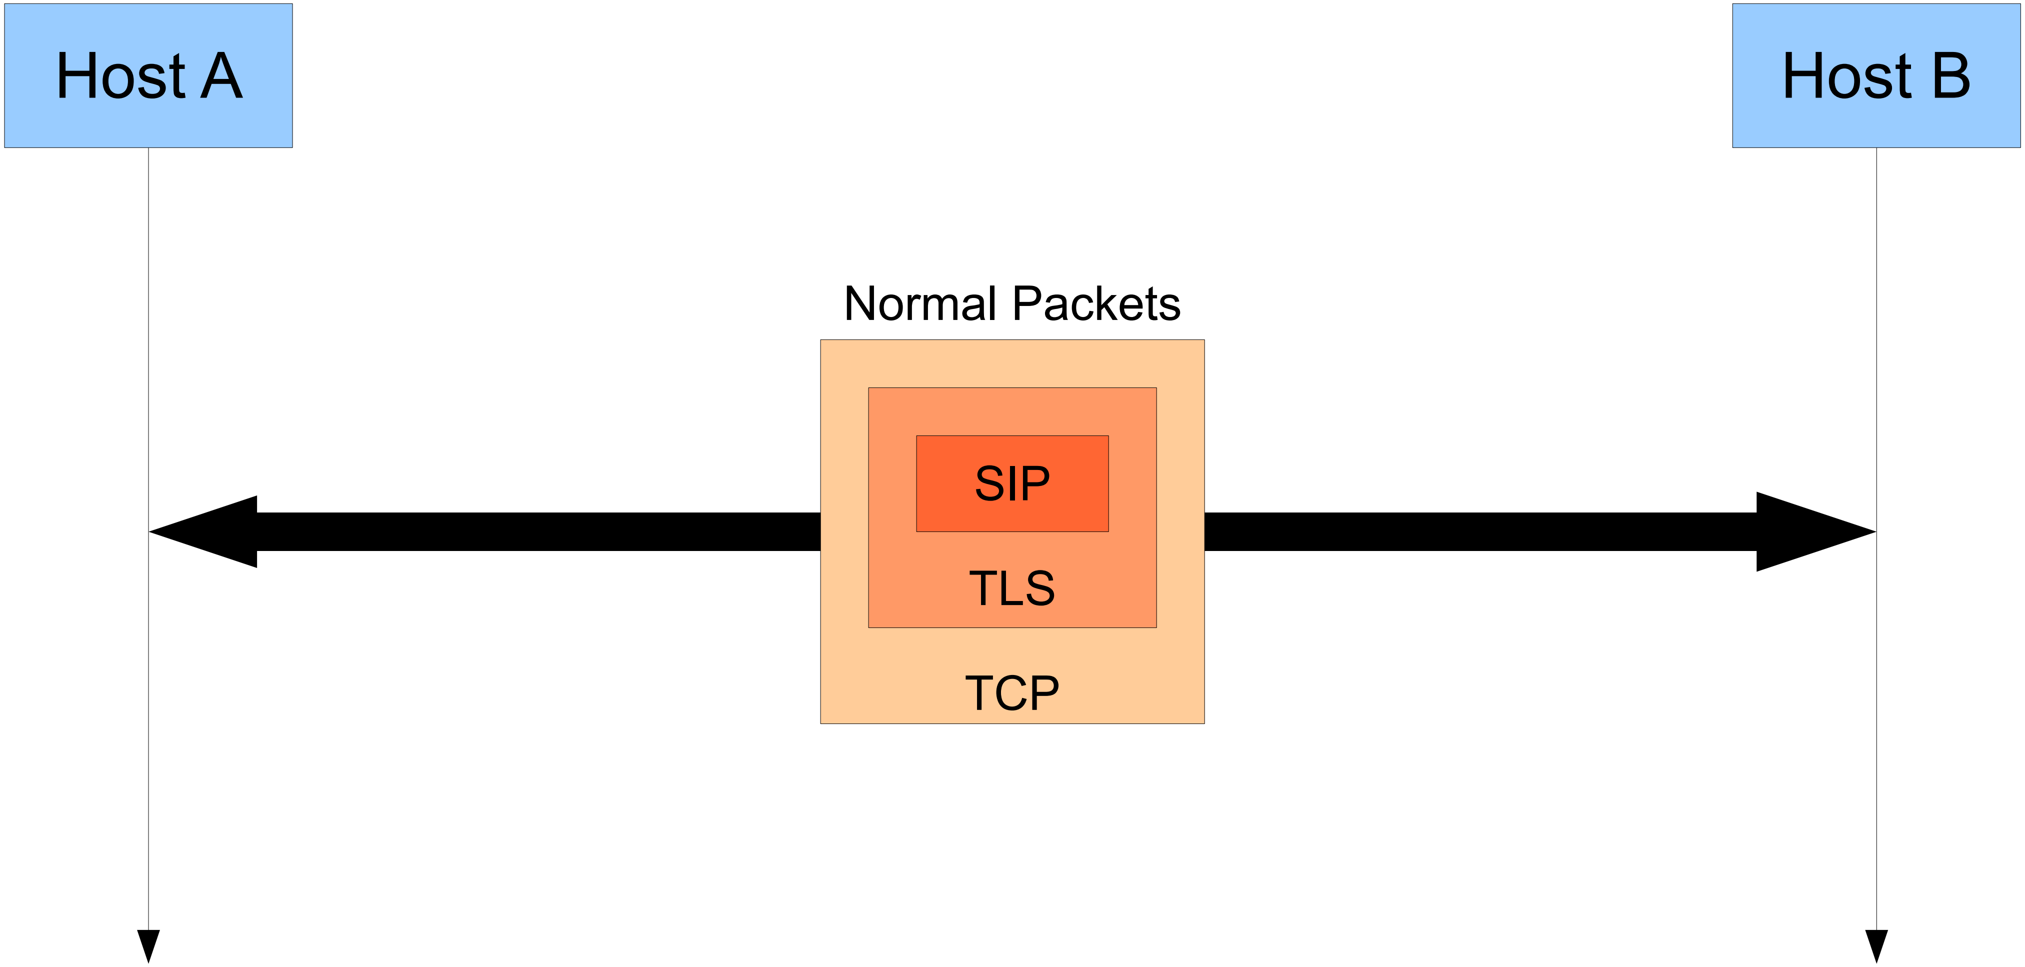

Concurrent Protocol Initialization

Whenever feasible:

- ***embed*** protocol-specific handshake info into graph
- ***run handshakes concurrently*** while negotiating

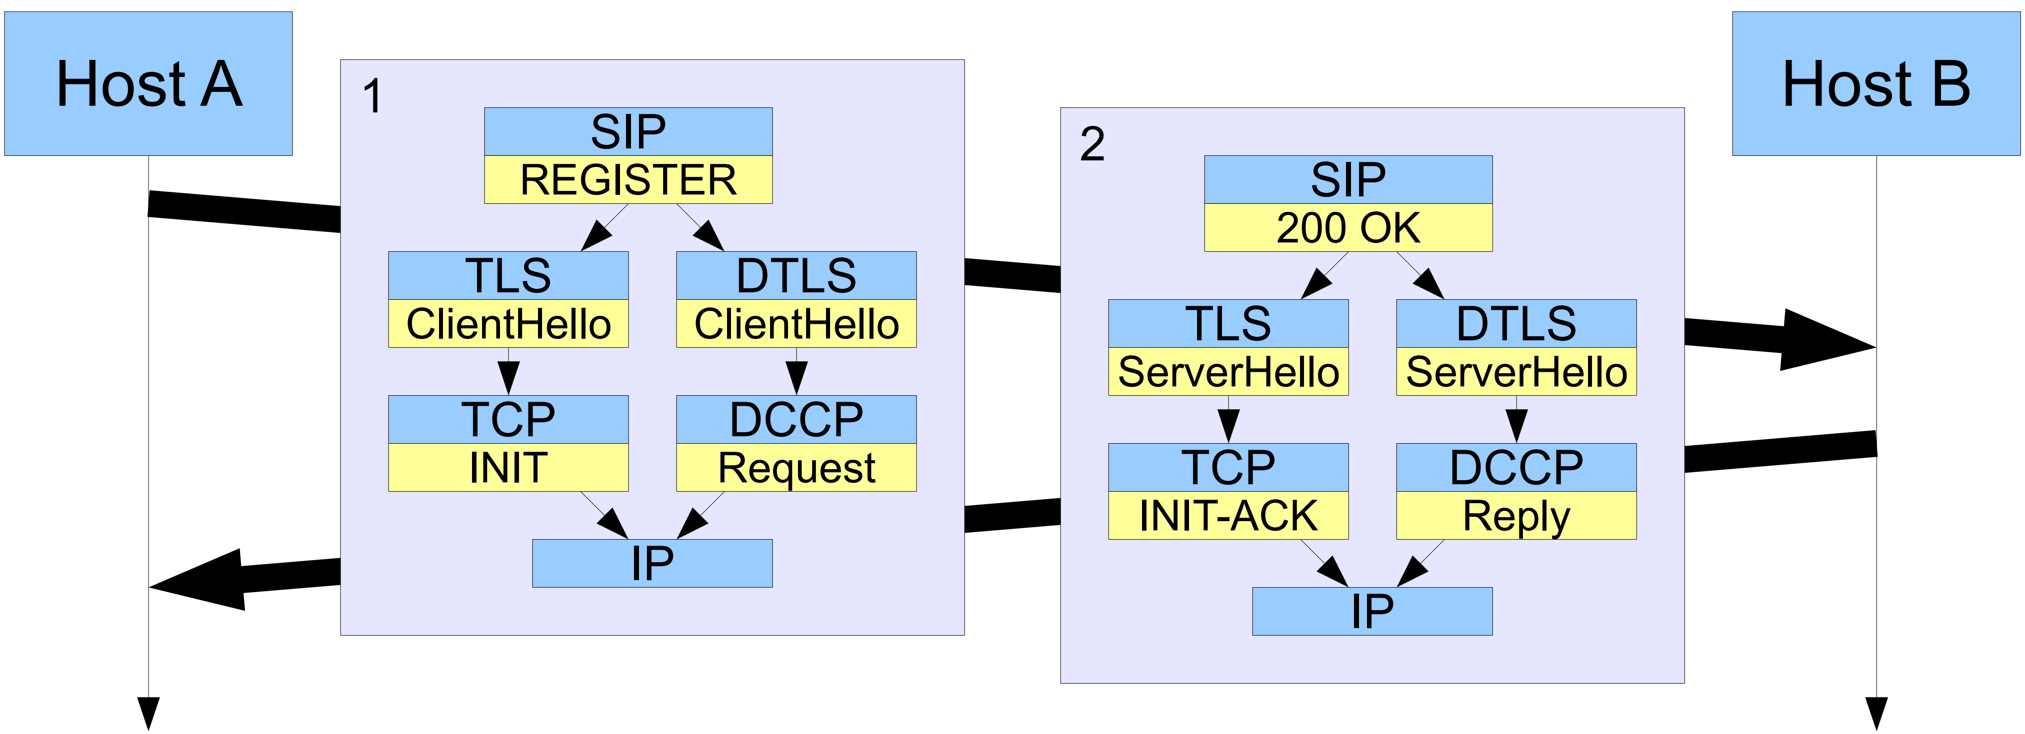

Key Benefits of Negotiation Model

- Happens strictly between nodes concerned
  - Users, Name server admins don't have to care
- Middleboxes can participate in process
- Protocol graph representation scales to handle:
  - Arbitrarily deep protocol stacks
  - Many alternatives per layer
- Setup whole "layer cakes" in minimal # of RTTs
  - With options

(For representing and transmitting graph, negotiation transport protocol, etc., see our HotNets '09 paper)

Contexts and Stacks

- ***Context*** ≡ underlying substrate; *cannot change*
- ***Stack*** ≡ protocols to be set up; *can change*

Example 1: Application-Level VoIP Protocol Stack Negotiation

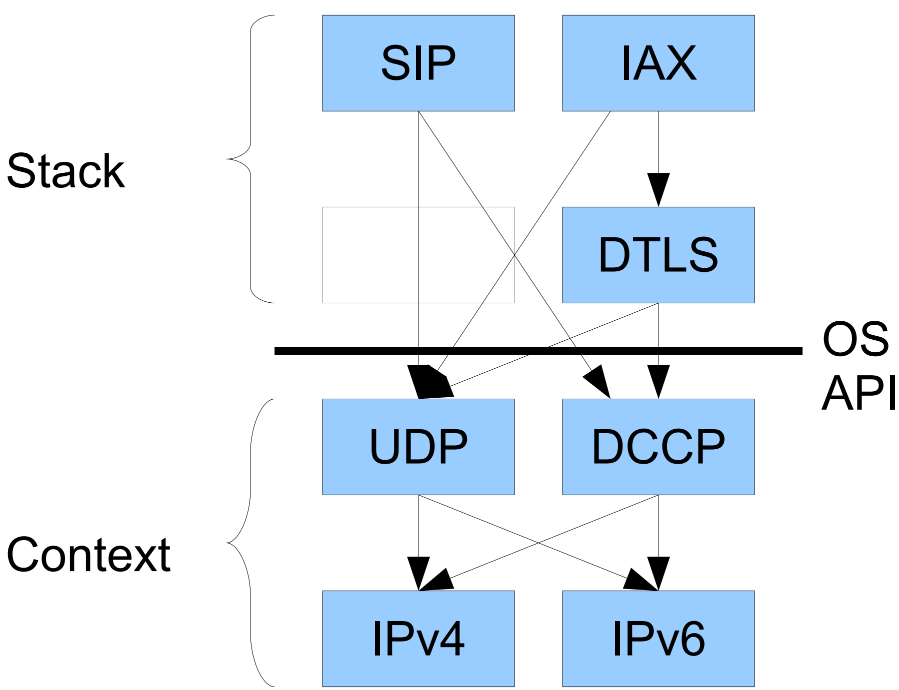

Example 2: OS-Level, Application-Transparent Transport Stack Negotiation

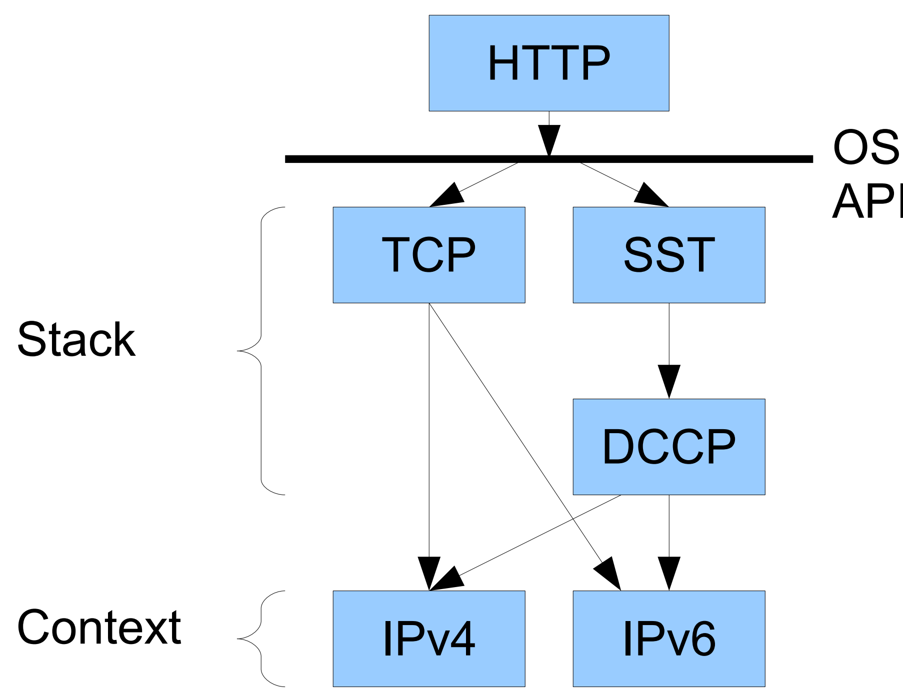

Negotiation Across Contexts

Scenario 1: Application-Level
VoIP Protocol Stack Negotiation

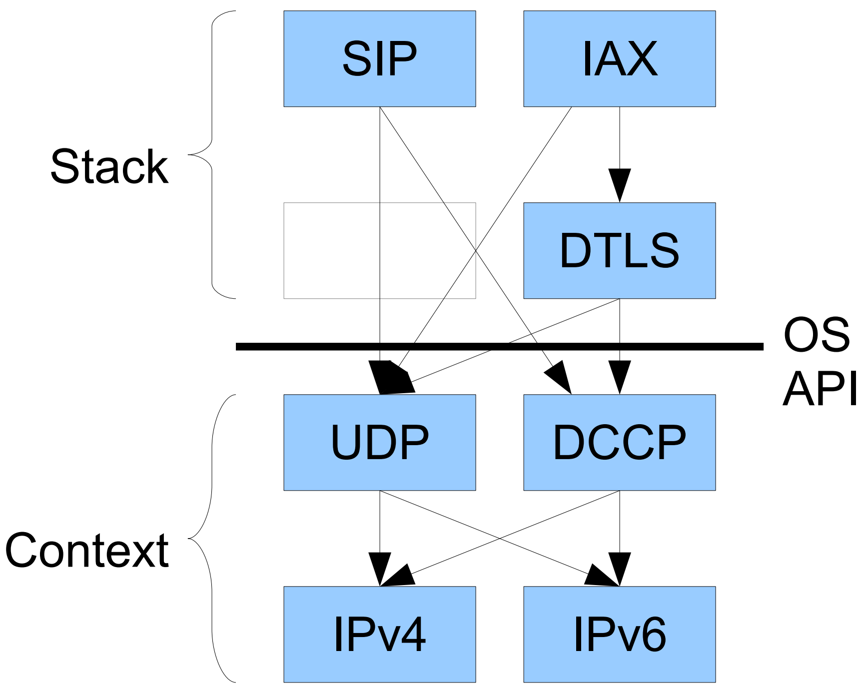

App can't send 1 packet
that's *both* UDP & DCCP!

Scenario 2: Application-Transparent
Transport Protocol Negotiation

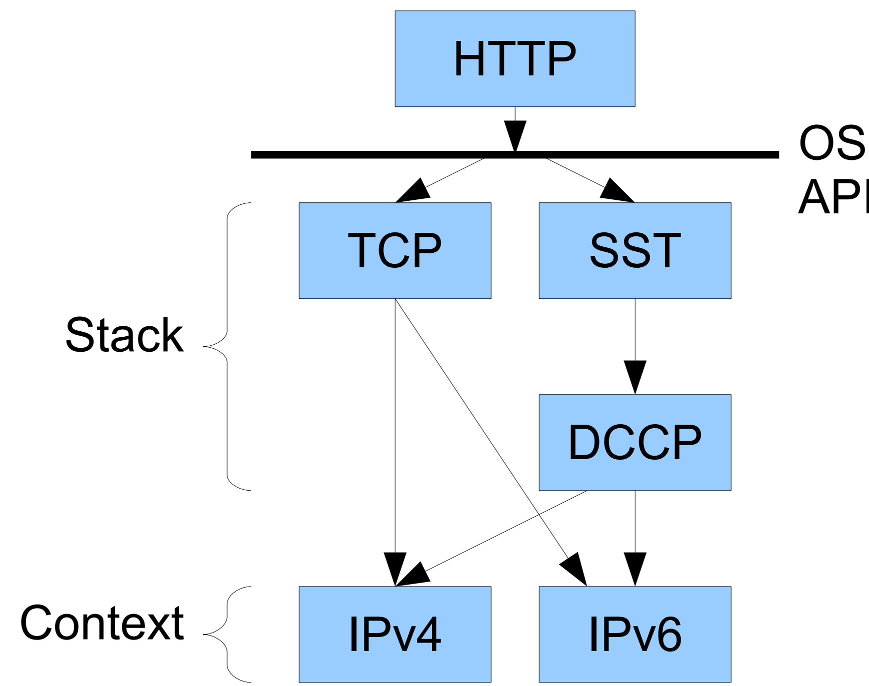

OS can't send 1 packet
that's *both* IPv4 & IPv6!

$\Longrightarrow$ *must* try each context separately

Combined Solution

1. Identify feasible communication **Context(s)**
   - e.g., UDP session (IP_a:$port_a$, IP_b:$port_b$)

2. Negotiate **Stack** within each context:
   a) Initiator sends a **Protocol Graph Proposal**
   b) Responder returns **Revised Protocol Graph**
   c) (Optional) further protocol graph revision steps
   d) Peers commit, **Acknowledge Protocol Graph**
   e) Communication proceeds via negotiated protocols

Combined Implicit/Explicit Solution

- Implicit, parallel negotiation *across contexts*
- Explicit, in-band negotiation *within a context*

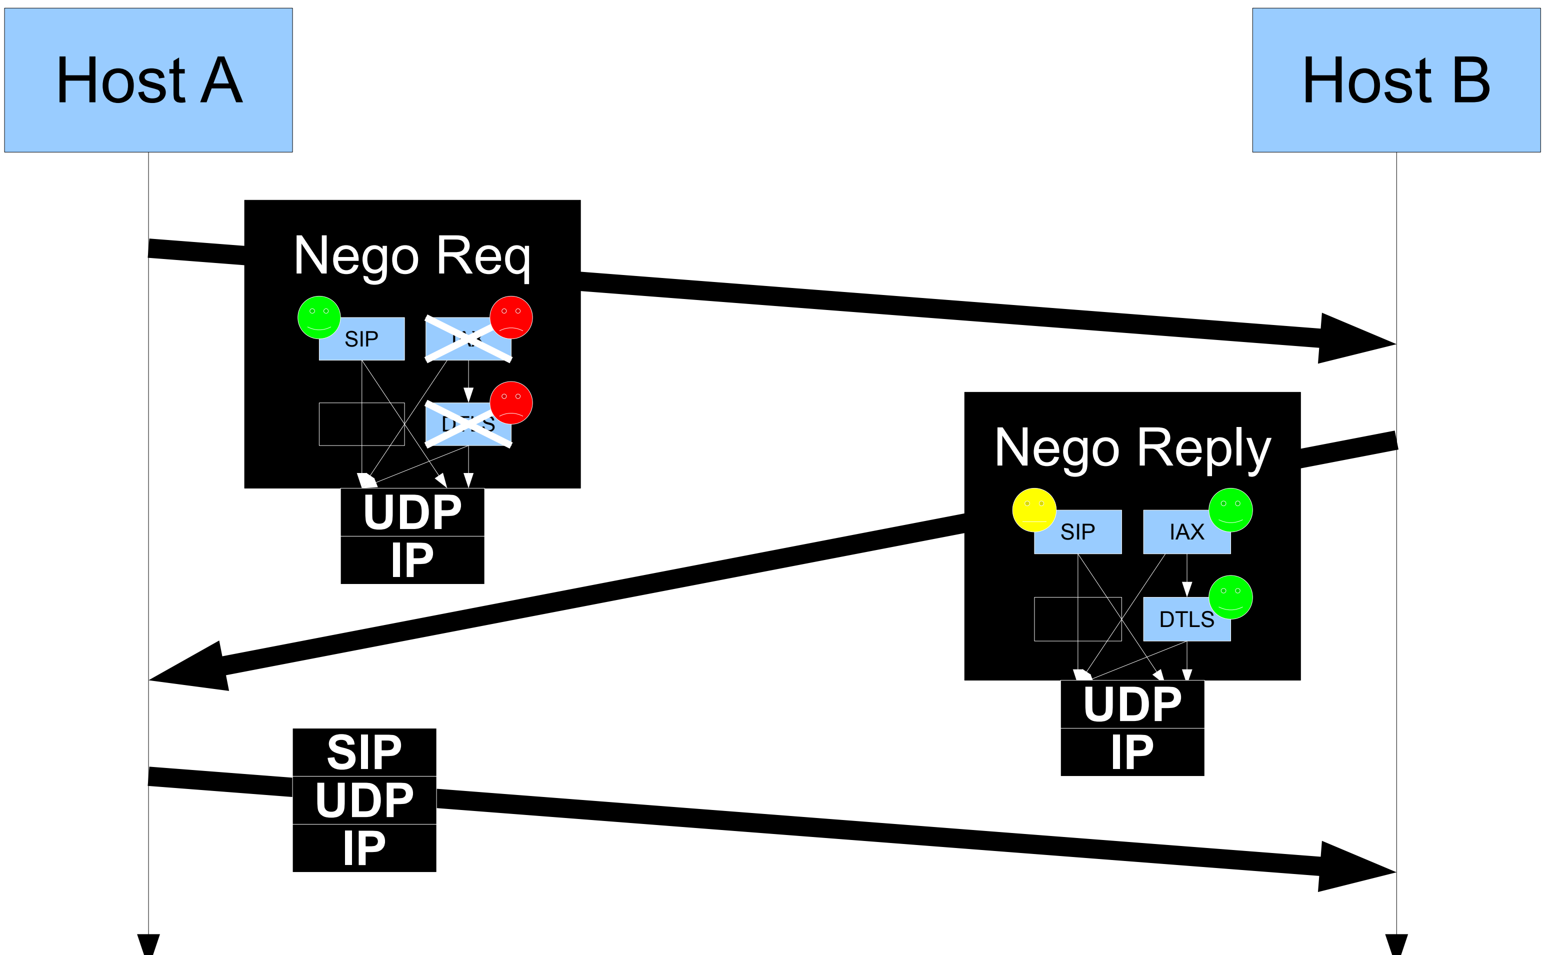

A Framework for Negotiation

Negotiation Strategies

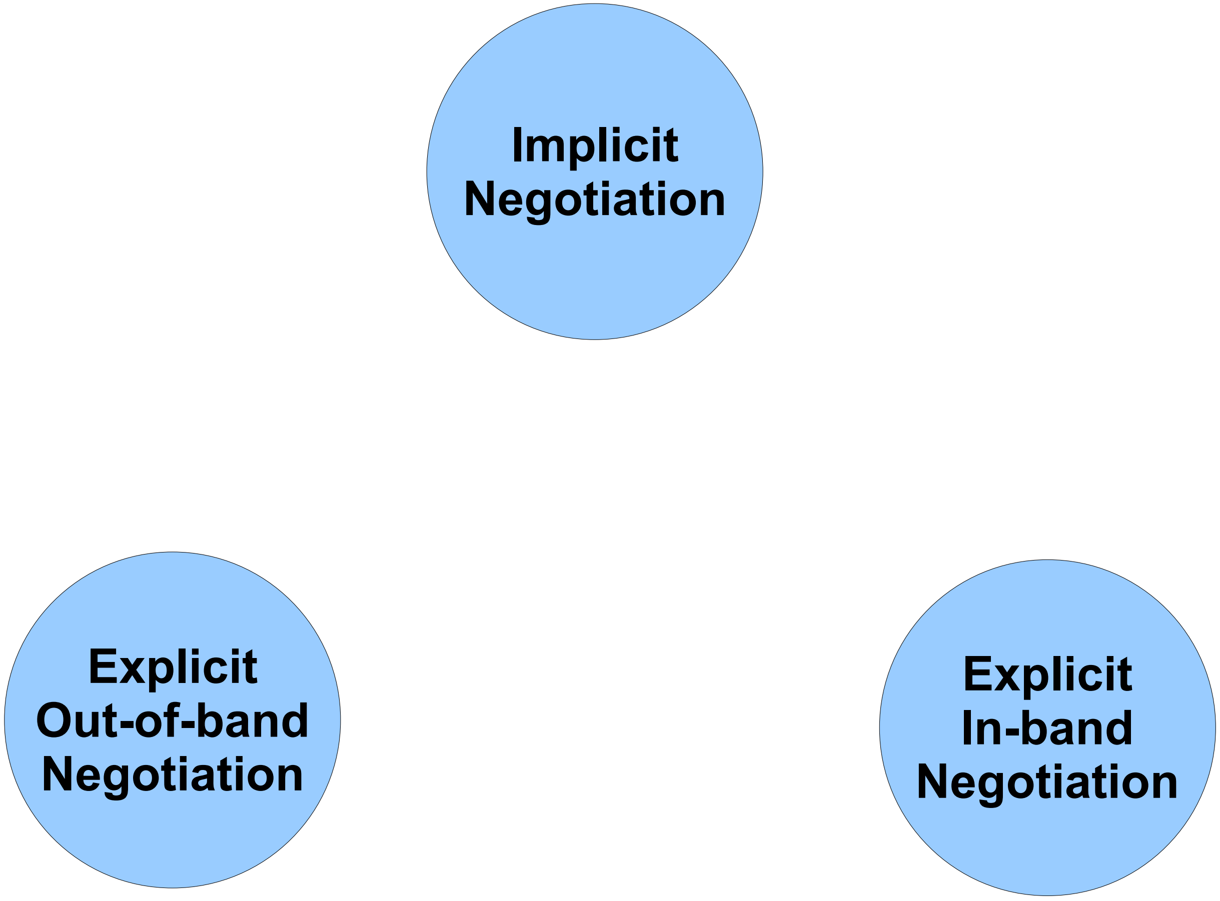

The Negotiation Triangle

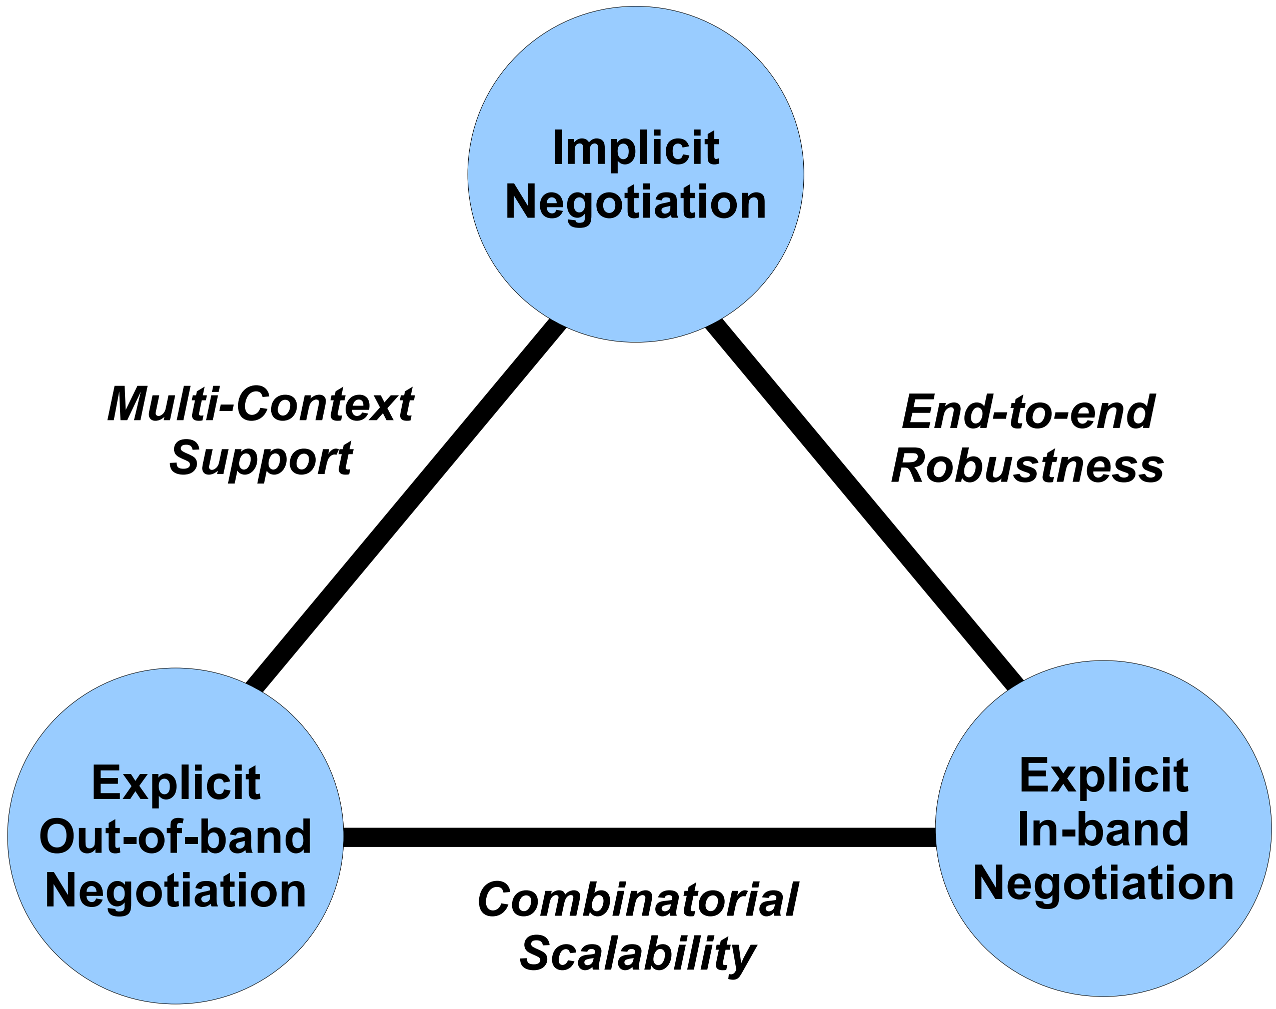

The Negotiation Triangle

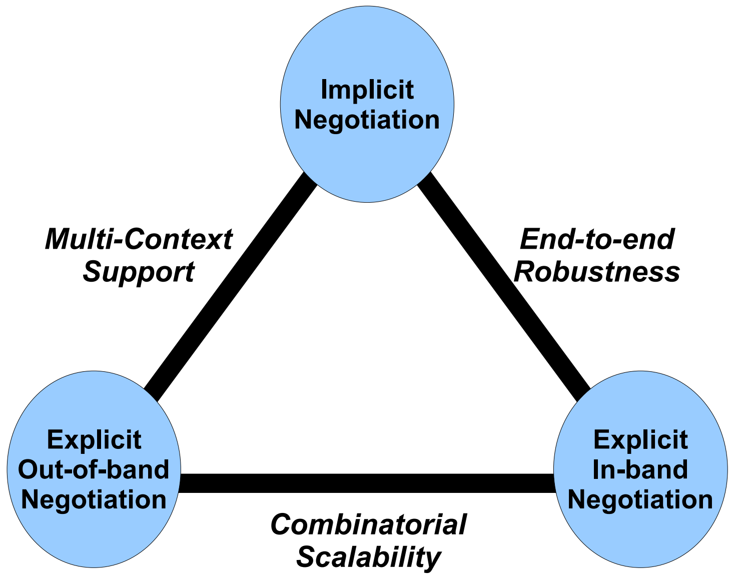

For any given negotiation strategy,
you get two of three desirable properties

To get all three properties,
a hybrid of at least two strategies is necessary

Arigato!

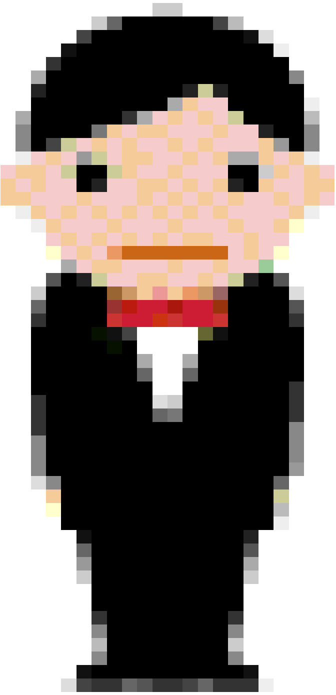

The floodgates are open!

(Please join tae@ietf.org for discussions)