

Introduction to Transport Network

Ivan Pepelnjak (ip@ipSpace.net)

ipSpace.net AG

The logo for ipSpace, featuring the text "ipSpace" in a white, cursive script font. The logo is positioned in the lower right quadrant of the slide, overlaid on a background of diagonal stripes in various shades of orange, yellow, and grey.

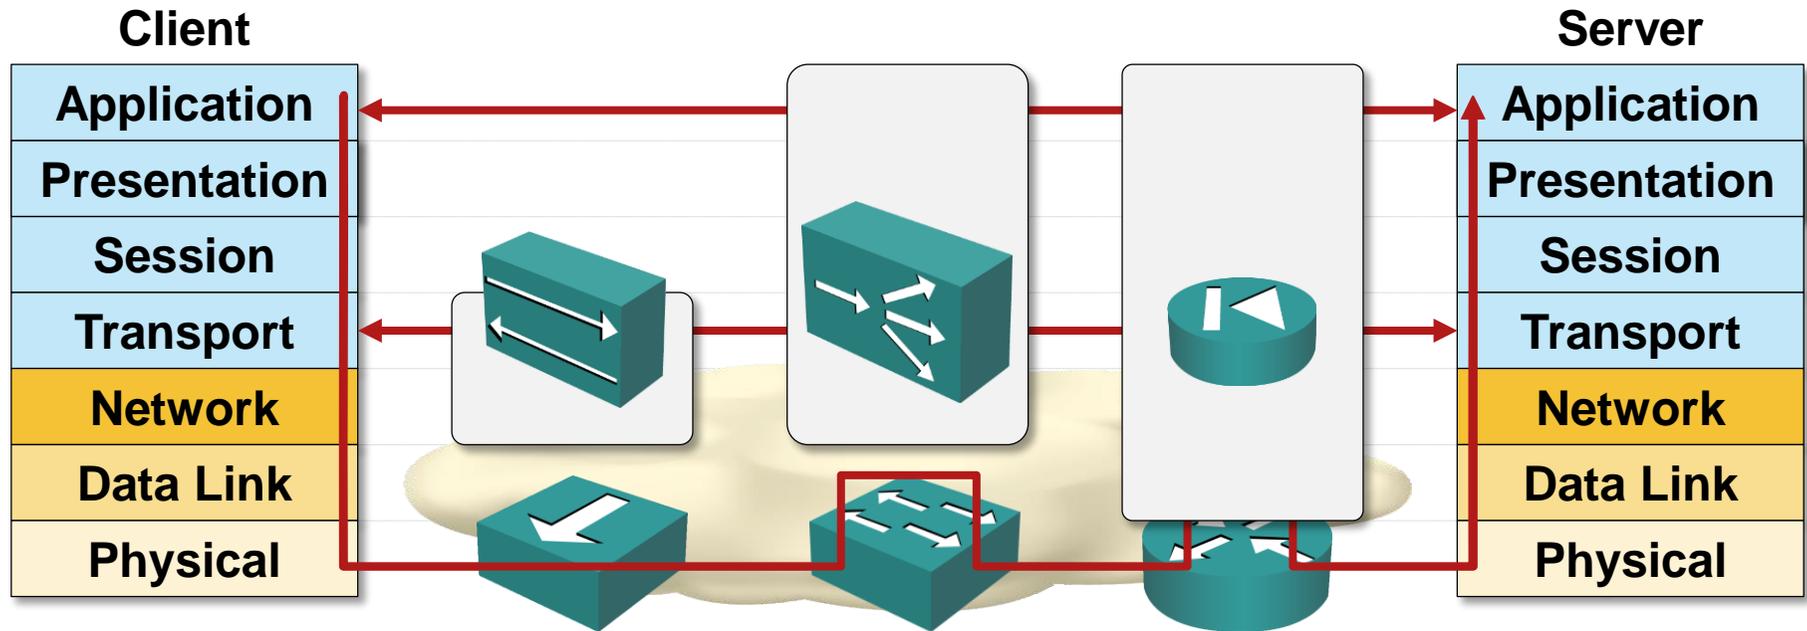
ipSpace

Focus Area: Transport



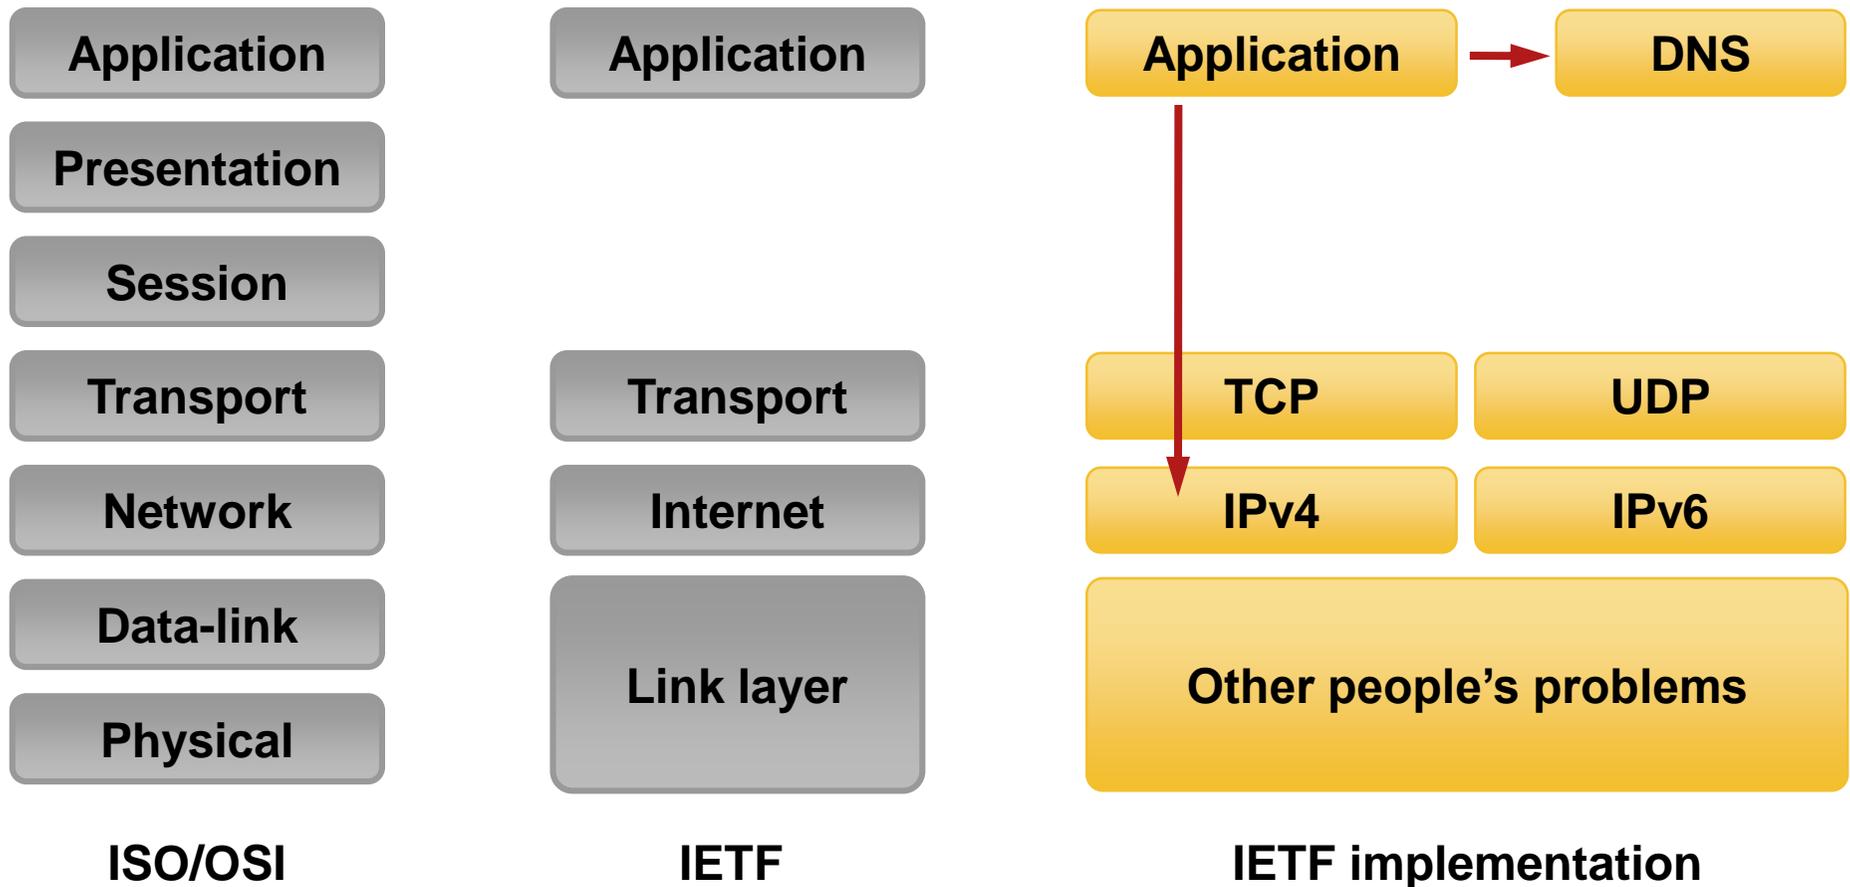
- What's going on behind the scenes?
- Fallacies of distributed computing
- TCP – basics, impact of latency/RTT and packet drops
- HTTP – principles, MIME types, CGI ...

Data Communications and Onions



U-Data	S-Port	D-Port	P	D-IP	S-IP	V	ET	VLAN	S-MAC	D-MAC
--------	--------	--------	---	------	------	---	----	------	-------	-------

TCP/IP: Broken Protocol Stack



- Session layer is missing (session endpoints tied to IP addresses)
- IP addresses are visible to applications (no L3-L7 abstraction)
- DNS is an optional add-on application

Broken Socket API

Ideal

```
conn = Network.Connect("example.com", "http")
```

TBD

OK

```
conn = new Socket("example.com", 80)
```

Java

Broken

```
memset(&hints, 0, sizeof(hints));
hints.ai_family = PF_UNSPEC;
hints.ai_socktype = SOCK_STREAM;
error = getaddrinfo("example.com", "http", &hints, &res0);
if (error) { errx(1, "%s", gai_strerror(error)); }

s = -1;
for (res = res0; res; res = res->ai_next) {
    s = socket(res->ai_family, res->ai_socktype, res->ai_protocol);
    if (s < 0) { cause = "socket"; continue; }

    if (connect(s, res->ai_addr, res->ai_addrlen) < 0) {
        cause = "connect";
        close(s);
        s = -1;
        continue;
    }

    break; /* okay we got one */
}

if (s < 0) { err(1, "%s", cause); }
```

Socket API

Consequences of Broken Stack & API

- Every application reinvents the wheel (sometimes badly)
- New network/transport protocols are a royal pain (see: IPv6) and require application changes
- Endpoint mobility is hard to implement (network-layer addresses have to move)
- No automatic transport/network protocol selection
- No load balancing across multiple endpoint addresses
- No automatic session reconnect

SCTP solves most of the above, but nobody is using it

Fallacies of Distributed Computing

1. The network is reliable.
2. Latency is zero.
3. Bandwidth is infinite.
4. The network is secure.
5. Topology doesn't change.
6. There is one administrator.
7. Transport cost is zero.
8. The network is homogeneous.

Source: Peter Deutsch, Sun Microsystems (1994), via Wikipedia

Always test your application in realistic environment

Application Testbed



Emulate realistic WAN transport

- Limited bandwidth
- High latency
- Packet drops
- Actual user devices (don't emulate iPhones) on actual connections

Tools: Traffic shaping/policing on routers/Linux, WANem
<http://packetlife.net/blog/2011/jan/12/emulating-wans-wanem/>

Questions?

