

Name based sockets

Javier Ubillos

Swedish Institute of
Computer Science

July 25th, 2010



<http://www.ietf.org/id/draft-ubillos-name-based-sockets-01.txt>

Name Based Sockets !

Making application development easier

The general problem

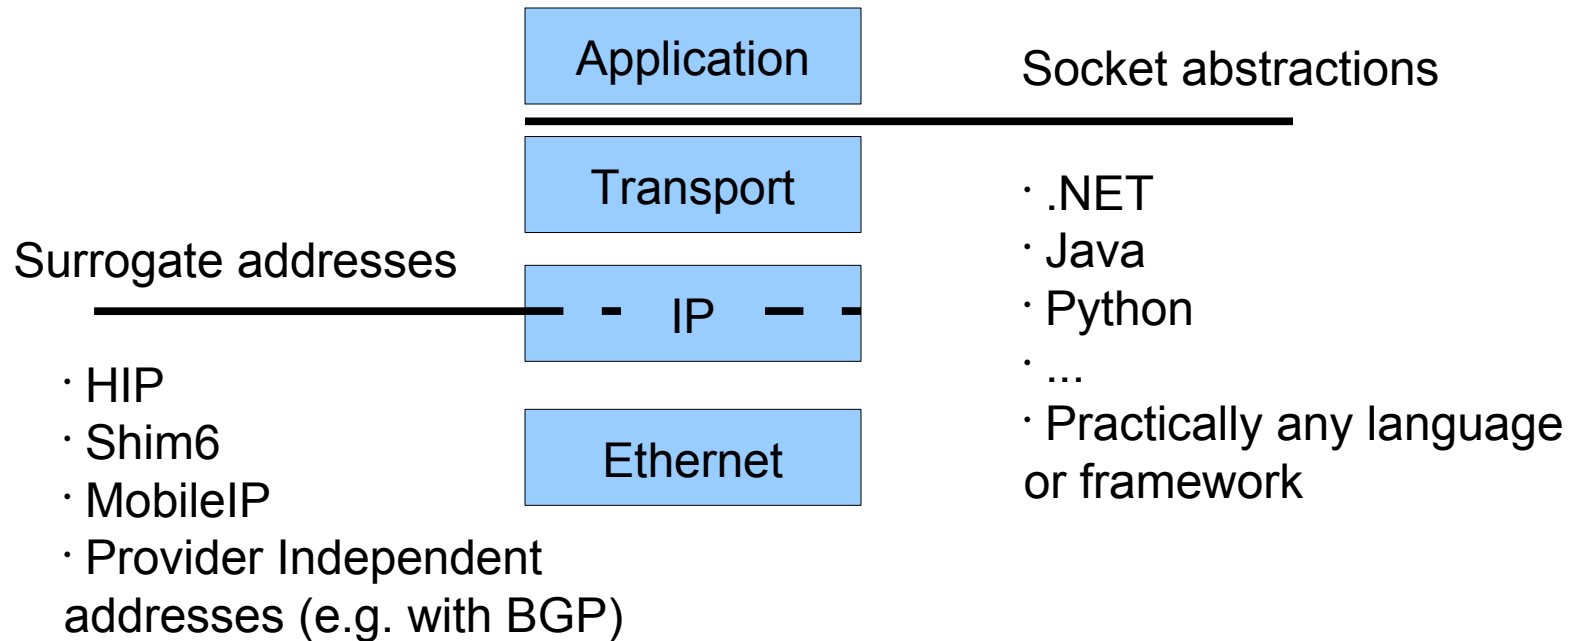
All IP (locator) management is done by the application.

Therefore, all interesting features need to be implemented by the application.

- Mobility
- Multi-homing
- IPv4/IPv6 interoperability
- NA(P)T traversal
- Path diversity exploitation
- Etc...

```
addr = gethostbyname( someString );  
...  
connect( ..., addr, ... );  
write( ... );  
close( ... );  
connect( ..., addr, ... );  
write( ... );  
close( ... );
```

Two typical approaches



Surrogate addresses

“Application transparency gives backwards compatibility (API)”

- Extra name spaces.
- Extra resolutions (more indirections)
- Applications are not aware, hence still might try to solve issues in app-space.

Surrogate addresses



The diagram on the right side of the slide shows a vertical stack of four light blue rectangular boxes representing protocol layers. From top to bottom, they are labeled: 'Applicati...', 'Transpo...', 'IP', and 'Ethere...'. A horizontal grey line is drawn across the stack, positioned between the 'Transpo...' and 'IP' layers. The text 'Surrogate addresses' is placed above this line, with a short horizontal line extending from the text to the main line.

Socket abstractions

Developers seem to like them...

- One implementation for every framework
- More often than not
 - Resolve once
 - Reuse IP
 - Reuse IP
 - Reuse IP
 - Reuse IP
 - Reuse IP
 - Reuse IP

Socket abstractions

Applicati

Transpo

IP

Etherne

What do we want?

- No new indirections
 - No new delays (e.g. first packet delay)
 - Address management
 - Mobility
 - Multi-homing
 - Renumbering
 - IPv4/IPv6 interoperability
 - NAT penetration
 - Backwards compatibility
-

Components

- API
- Initial name exchange
- Address updates
- Backwards compatibility (on the road map)

API


```
fd = socket( AF_NAME, SOCK_STREAM, 0);  
struct sockaddr_name name_sock;  
  
// Initialize name_sock with remote name  
  
bind( fd, name_sock, sizeof(name_sock));  
connect( fd, name_sock, sizeof(name_sock));  
  
write(fd, send_buffer, len);  
read(fd, recv_buffer, len);
```

The components (API)

- listen() - Prep for incoming session
 - _ fd = listen(local_name, peer_name, service, transport);
- open() - Initiate outgoing session
 - _ fd = open(local_name, peer_name, service, transport);
- accept() - Receive incoming session
 - _ accept(peer_name, fd);
- read() - Receive data
 - _ data = read(fd);
- write() - Send data
 - _ write(fd, data);
- close() - Close session
 - _ close(fd);

Initial name exchange

Traditional




Name: `host.left.org`
IP: `a.b.c.d`

Application

Operating
system


Name: `host.right.org`
IP: `w.x.y.z`



Operating
system

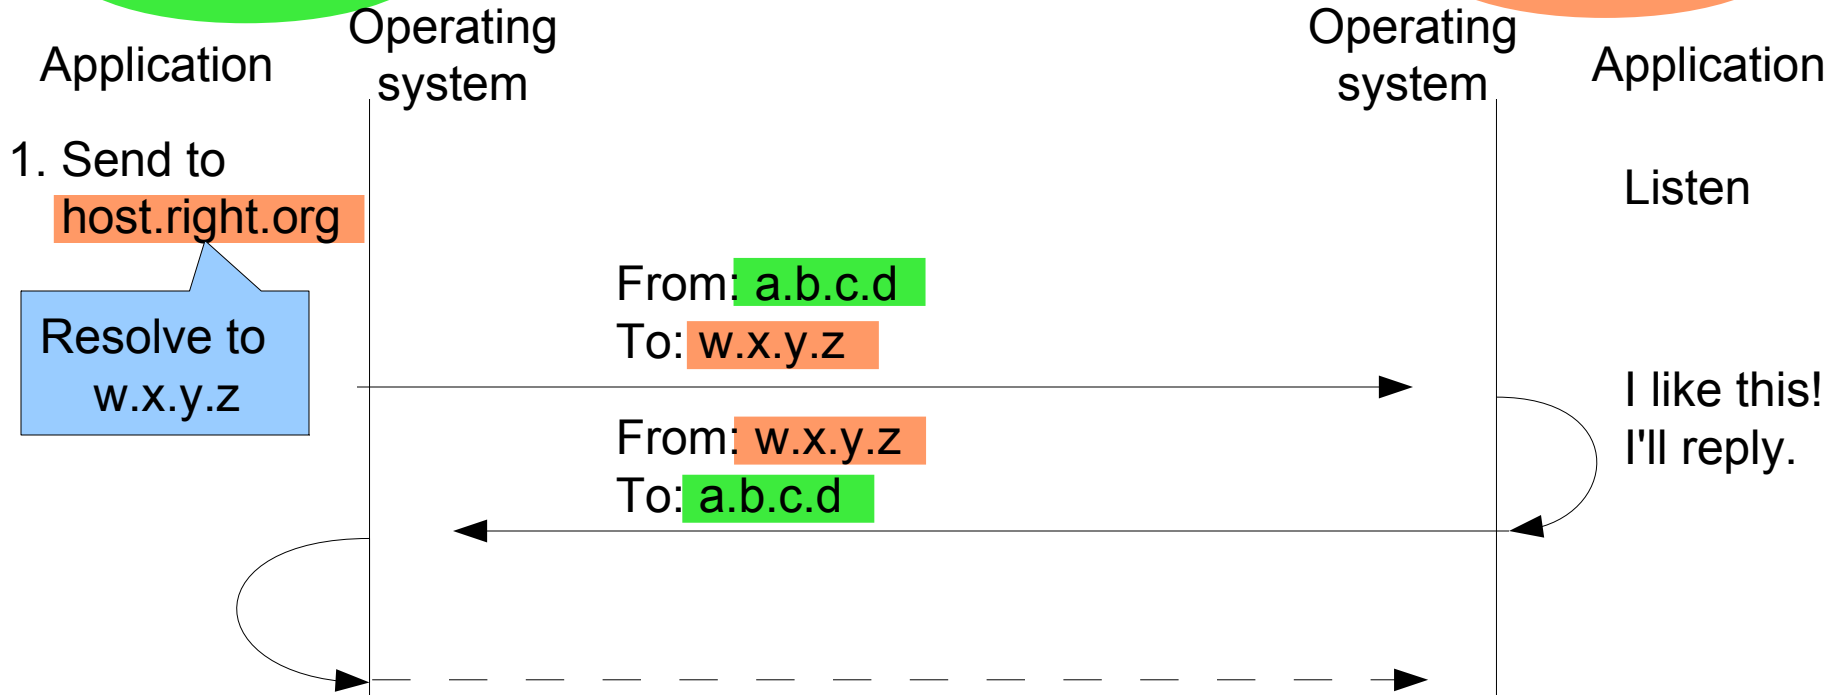

Application

Traditional

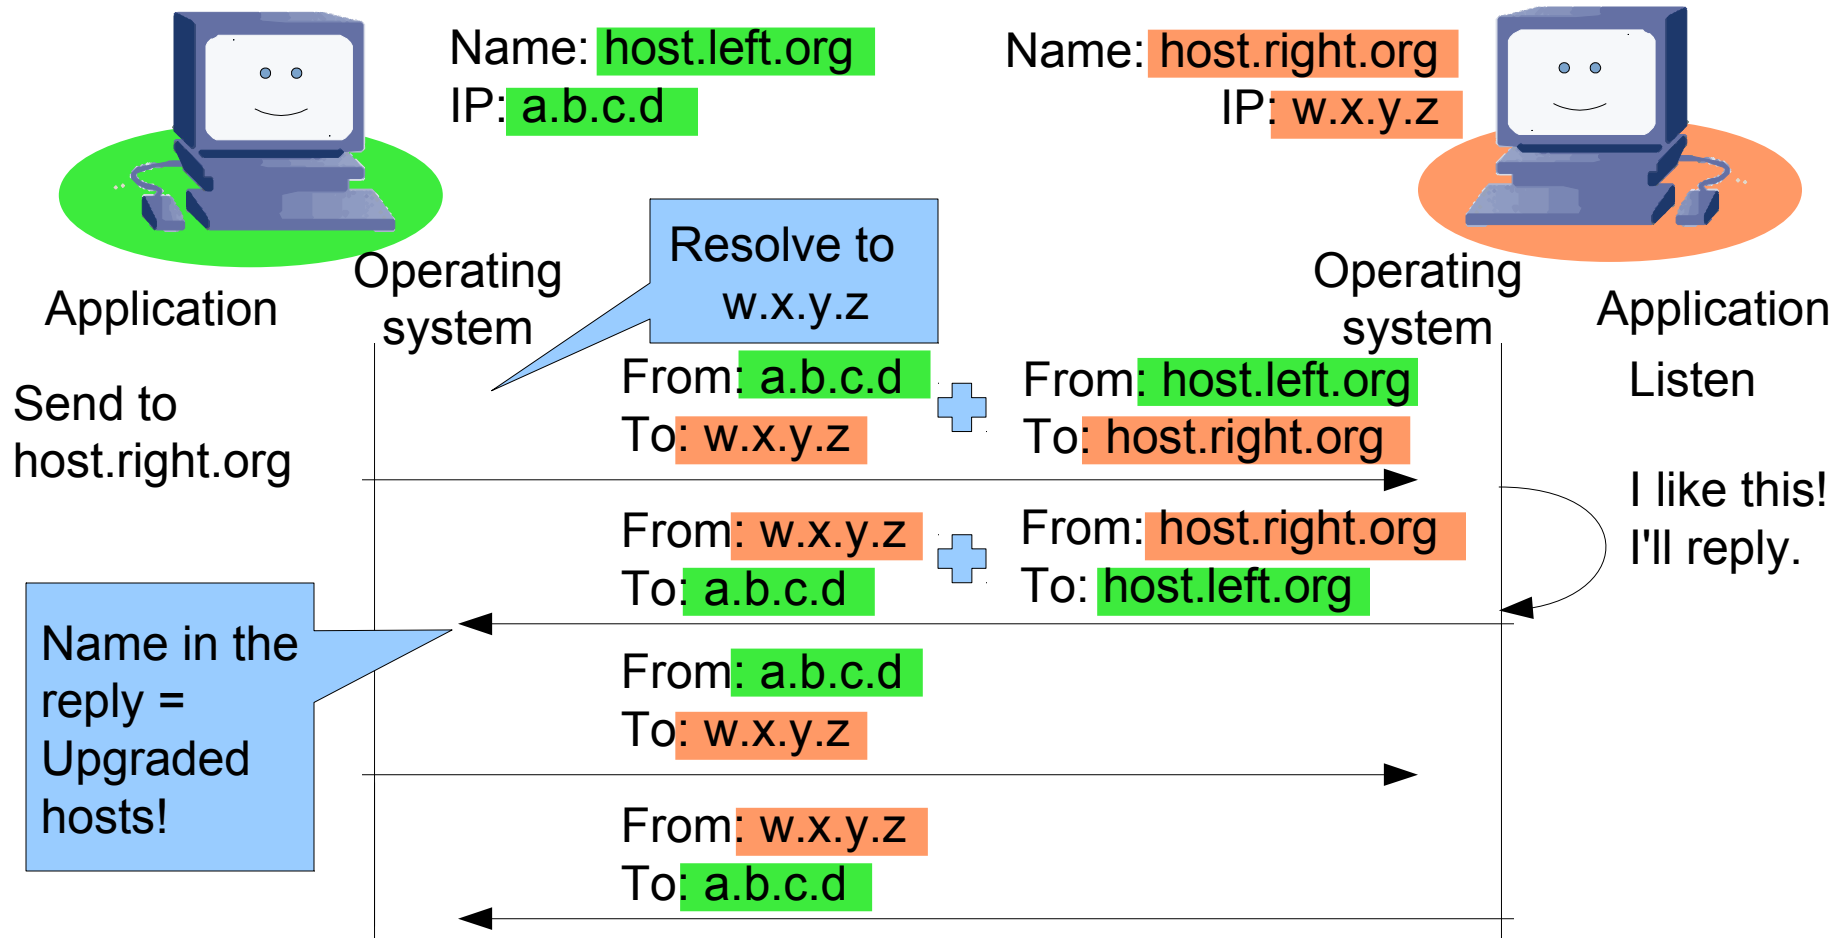


Name: **host.left.org**
IP: **a.b.c.d**

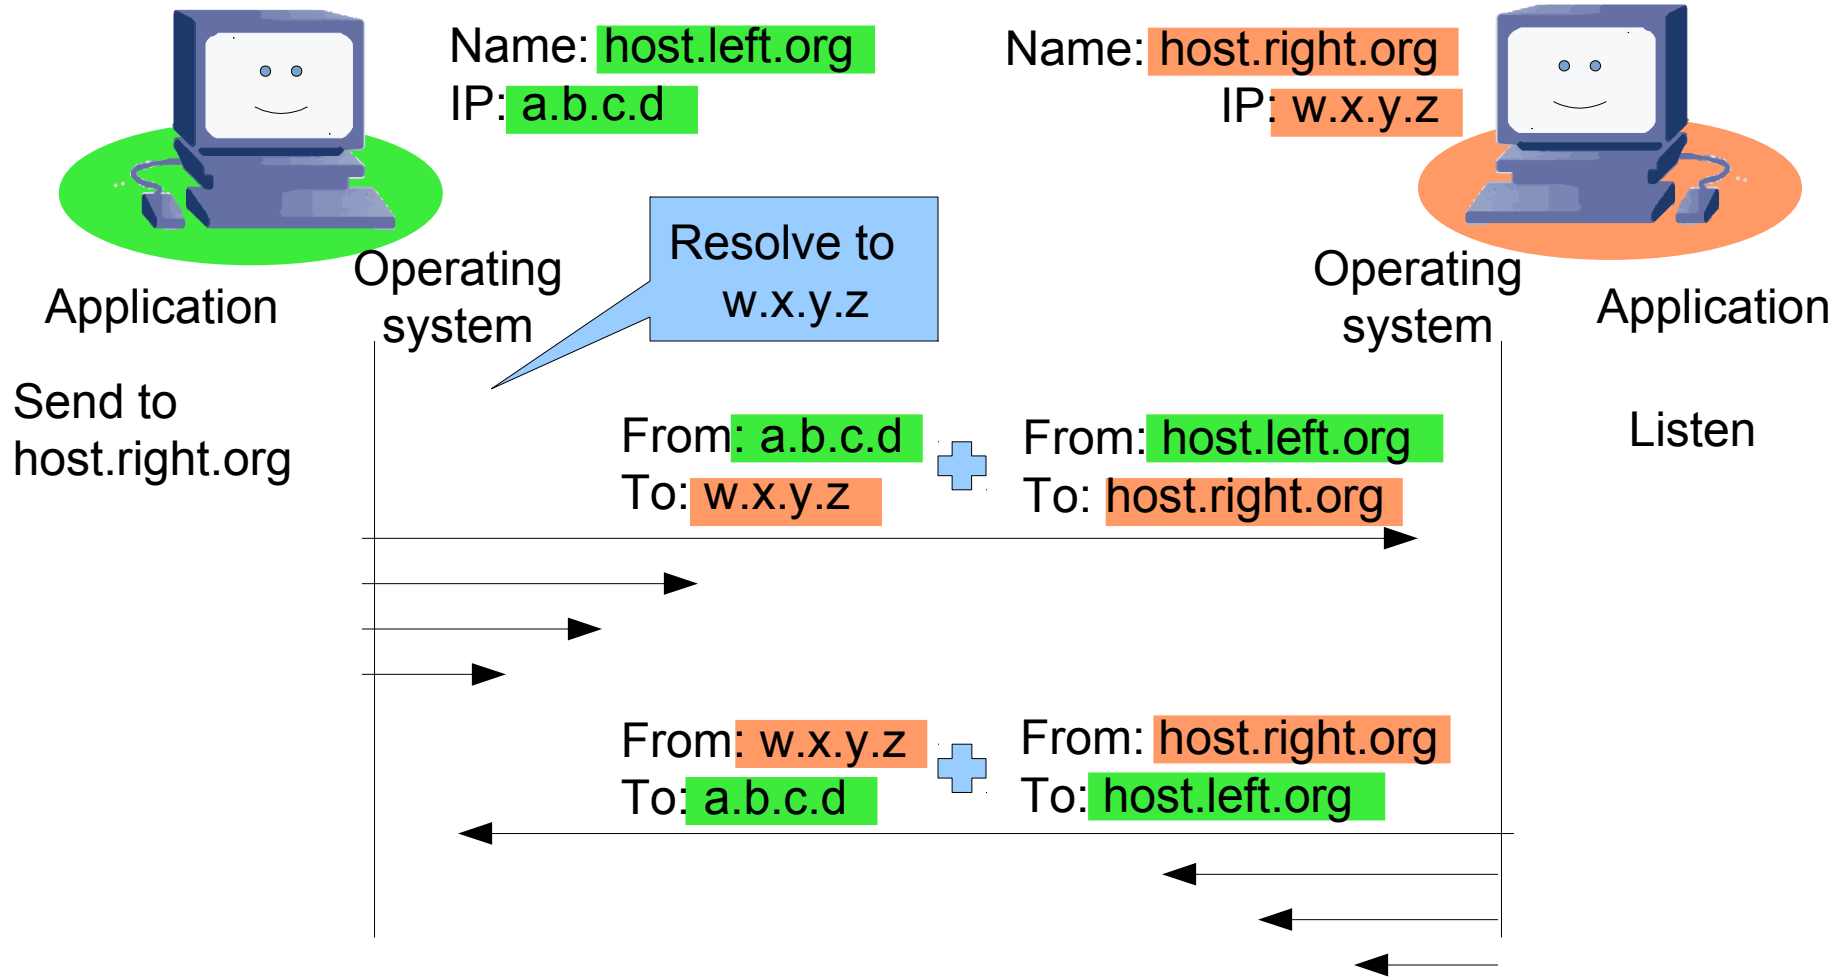
Name: **host.right.org**
IP: **w.x.y.z**



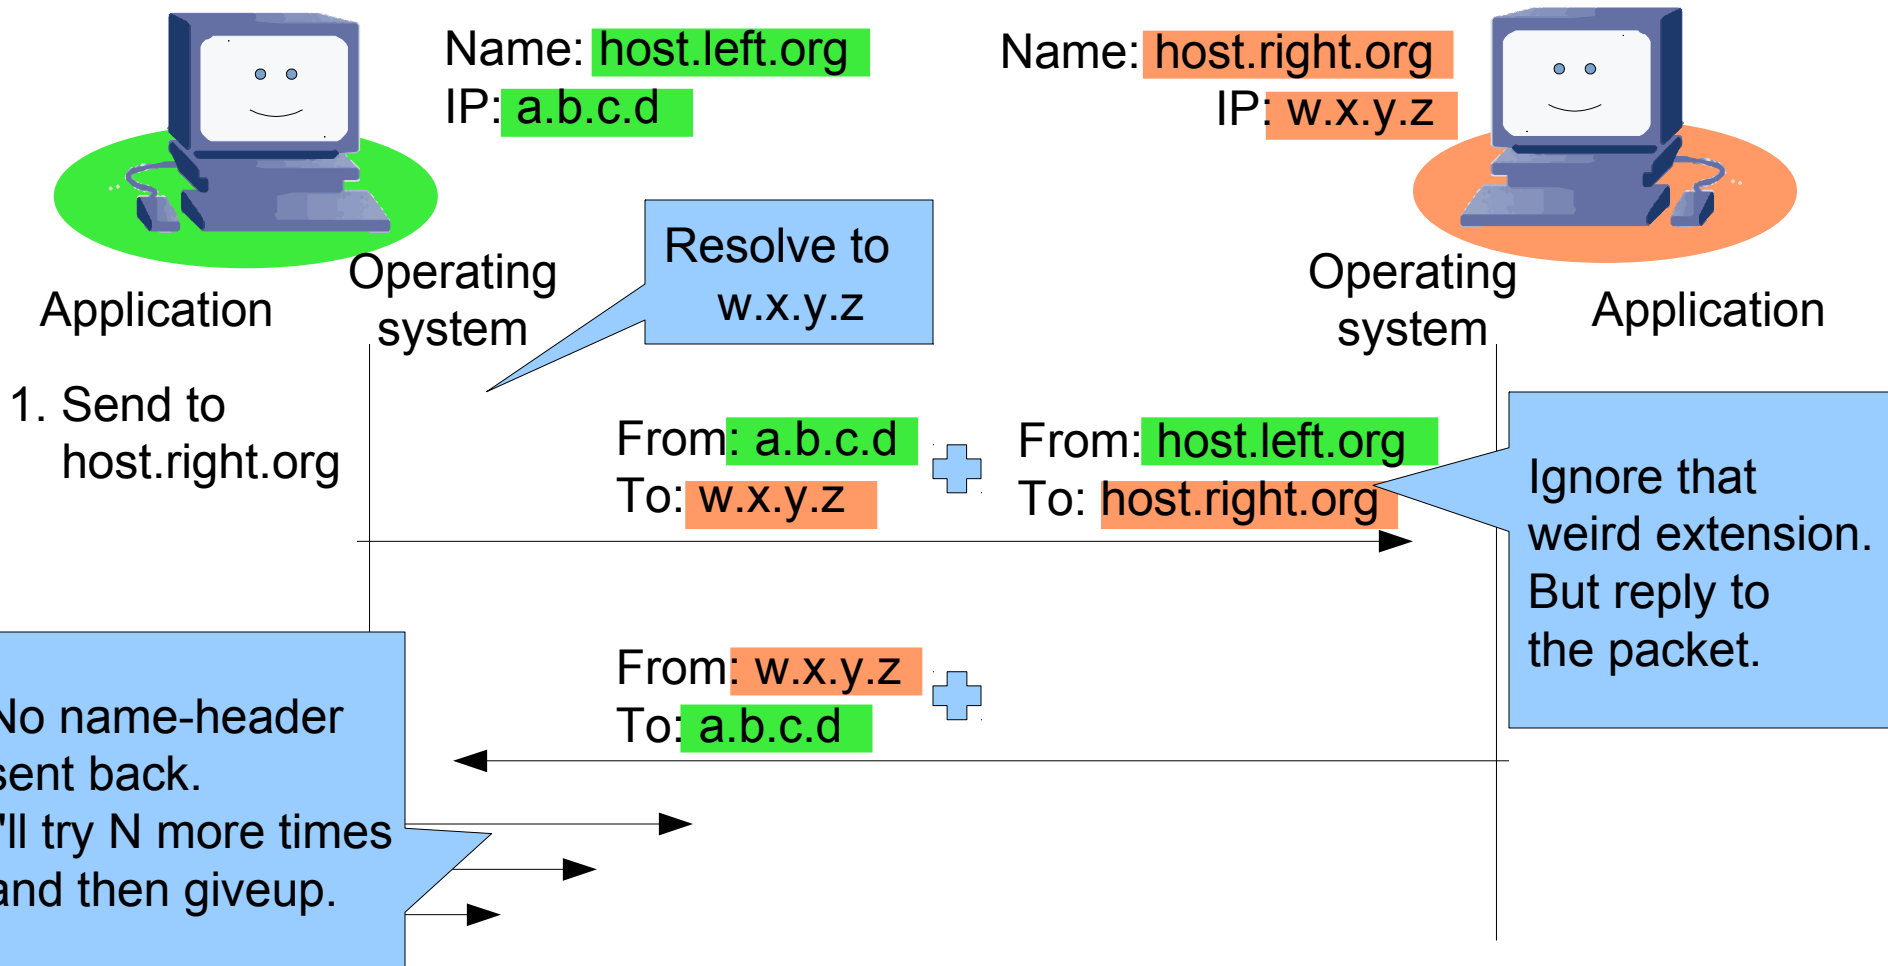
Name exchange



Name exchange



Backwards compatibility



The current prototype

- Supports TCP
 - _ Uses TCP semantics
 - socket(), listen(), open(), accept(), read(), write()
- Supports Shim6
 - _ Well, to a certain extent, we are working on it :)
- Exchanges names
- Linux
 - _ Ubuntu (client/server)
 - _ Android (client)

Implementation by Juan Lang (UC Davis)
and by Zhongxing Ming (Tsinghua University)

UC DAVIS
UNIVERSITY OF CALIFORNIA



Current development

- Support for UDP
 - Using TCP-like semantics
- Mobility/Multi-homing
 - Shim6
- Collaboration between
 - Ericsson
 - Tsinghua University
 - Swedish Institute of Computer Science



The road map

- IPv4/IPv6 Interoperability
- NAT penetration
- Path diversity utilization
- Naming resolution (depth)
 - Host
 - Application
 - Etc...
- And more... Do you have any suggestions?
Please let us know!

Questions?