SunSoft
A Sun Microsystems Company

NFS and IPv6

SUN Microsystem Inc.

Sumandra Majee

smajee@eng.sun.com

SunSoft
A Sun Microsystems Company

Outline

- Overview on TI-RPC
- Design Goals
- Implementation
- Conclusion

SunSoft
A Sun Microsystems Company

Overview on TI-RPC

Client Side API

Client = clnt_create(host, prog, vers, nettype)

- Select network protocol depending on nettype
- Find remote host address and port
- Open communication channel

SunSoft
A Sun Microsystems Company

Overview of TI-RPC

Server side API

Server handle = svc_create(dispatch, prog, vers, nettype)

- Select network protocol depending on nettype
- Register service with portmapper and rpcbind
- Open communication channel

SunSoft
A Sun Microsystems Company

Overview on TI-RPC

What is in nettype

nettype is selected via /etc/netconfig

netid	semantic	flag	prot_family	prot	device
udp	tpi_clts	v	inet	udp	/dev/udp
udp6	tpi_clts	v	net	udp	/dev/udp6

- New Devices /dev/tcp6, /dev/udp6

SunSoft
A Sun Microsystems Company

Design Goals

- Backward Compatibility
- Easy interoparability between IPv4 host and dual stack IPv4/IPv6 hosts
- Minimal performance degradation for IPv4 path.
- Easy migration of TI-RPC application

SunSoft
A Sun Microsystems Company

Implementation

Dual Stack of IPv4 and IPv6

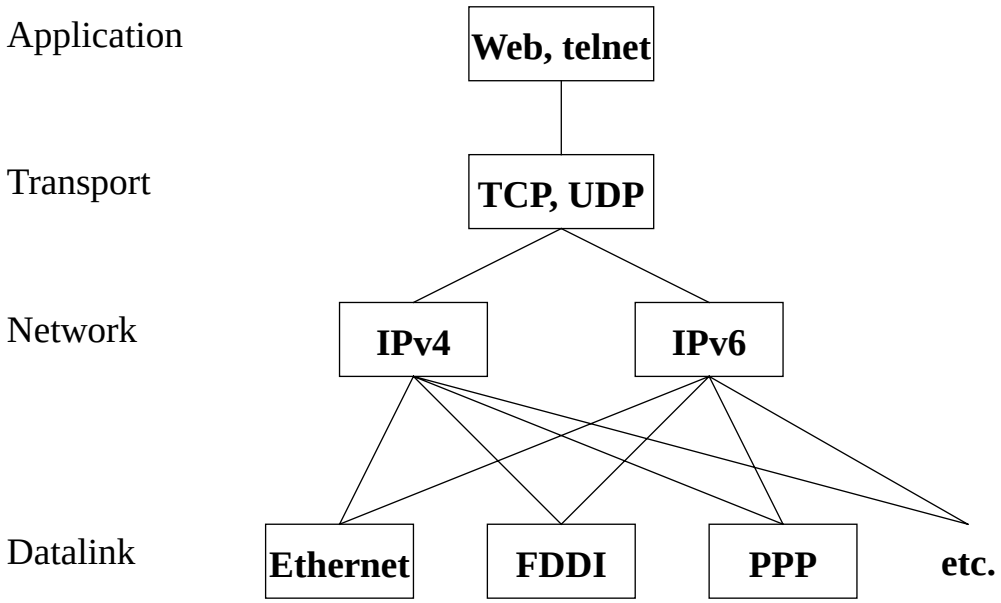

SunSoft
A Sun Microsystems Company

Implementation

Server

1. open an endpoint for communication
2. bind the server address and port
3. *register service with both IPv6 and IPv4*
4. listen for new request

Client

1. open an endpoint for communication
2. find remote host address
3. *contact rpcbind/portmapper with IPv6 or IPv4 first. If fails use the other protocol. This order can be customized via /etc./netconfig.*
4. Bind the server address and port
5. invoke clnt_call to access the remote procedures.


SunSoft
A Sun Microsystems Company

SunSoft
A Sun Microsystems Company

Implementation

- No Broadcast. RPC broadcast is achieved by subscribing to well known multicast address.
- Can not use IPv6 address (: notation) for most of the commands e.g mount.
- New netids for IPv6 (tcp6 and udp6)
- Universal address of IPv6 host is IPv6addr.porthi.portlo
- rpcbind database uses new netid (tcp6, udp6) for IPv6 based services.
- Many changes to reflect IPv6 address (sockaddr_in6)

SunSoft
A Sun Microsystems Company

Conclusion

- Designed for easy migration
- Minimize performance impact on IPv4 based NFS/RPC
- Preliminery results show about ~10% throughput degradation.