

Software that sends data across a network must understand...

- how to avoid collisions between packets,
- how to convert digital data to analog signals,
- how to detect and correct errors,
- how to route packets from one host to another,
- and much more.

What Happens When You Run a Browser to Retrieve a Webpage

- The browser in the application layer sends a request to the transport layer on your local machine.
- The transport layer breaks up the request into TCP segments, adds some information, and passes the request to the local internet layer.
- The local internet layer fragments the segments into IP datagrams and passes them to the host-to-network layer for actual transmission.
- The host-to-network layer encodes the digital data as analog signals appropriate for the physical medium and sends the request out the wire.

- The host-to-internet layer on the remote system receives the signals, decodes the analog signals back into digital data, and passes the IP datagrams up to the internet layer.
- The internet layer does some checking, reassembles any fragmented datagrams, and passes them to the transport layer.
- The transport layer checks that all the data arrived, requests retransmission of any missing or corrupt pieces, reassembles the packets into the original form, and passes the request to the application layer and the web server on the remote machine.
- The remote machine responds to the request, and the whole process is carried out in reverse.

Client-Server Paradigm or Model

- client - application software that initiates communication.
- server - application software that passively waits to be contacted.

Note: Client and server are SOFTWARE PROGRAMS and NOT the computers on which they run.

Multiple Services on One Computer

Requires:

- a fast processor and large memory.
- a multi-tasking OS (such as UNIX) that allows several application programs to run concurrently.

Advantages

- Hardware resources are shared by multiple services.
- The server computer needs only one physical connection to the network.
- Management overhead is reduced with only one computer to maintain.
- Server demand is often sporadic, and so an idle server does not use a computer's CPU while waiting for a request to arrive.

Disadvantages

- One server can affect others by exhausting or monopolizing computer resources.
- A computer failure can bring down multiple servers.

Identifying a Service

Question: With several servers on the same machine, how are they distinguished from one another?

Answer: Use a unique identifier.

- TCP and UDP use a protocol port number or simply a port.

Multiple Copies of a Server for the Same Service

Question: What happens when several clients request the same service at the same time?

Answer: Run multiple copies of the server concurrently.

Server programs typically consist of two parts:

- Master server (main thread server)
- Slave server (service thread)

Keeping Things Straight

Question: If there are multiple copies of a server, how does a client interact with the correct copy?

Answer: Use unique identifiers.

With TCP,

- client: (IP address, source port number)
- server: (IP address, destination port number)