

FTP, SMTP, DNS

1. The file transfer protocol (FTP) is one of the earliest Internet protocols. It is used to copy files between servers, and delete files on remote servers. It has a persistent control connection, on which a user issues commands and the remote site sends status responses. For each file transferred, it has a new data connection. The TCP port number of the file connection is sent over the control connection. What are the benefits and disadvantages of using a separate control channel?
 - benefits:
 - you can initiate multiple transfers, terminate transfers, etc. while the transfer is happening, without having to wait for transfer to complete
 - maintenance of state e.g. current directory
 - designed in days of unreliable control connection: signal client-server connection is healthy
 - control bandwidth
 - disadvantages:
 - additional overhead associated with maintaining multiple TCP connections
2. Describe the commonalities and differences between SMTP and (persistent) HTTP.
 - commonalities:
 - can persistent TCP in transport layer
 - both transfer files from host to client
 - differences:
 - HTTP operates between server and client (typically browser), while SMTP operates between mail servers
 - HTTP 80 vs SMTP 25
 - HTTP primarily pull; SMTP is a push protocol
 - SMTP messages must be ASCII, HTTP messages don't have this limitation
 - HTTP: one response per object; SMTP: all objects in one response
3. Suppose that John just set up an auto-forwarding mechanism on his work email address, which receives all of his business-related emails, to forward them to his personal email address, which he shares with his wife. John's wife was unaware of this, and activated a vacation agent on their personal account. Because John forwarded his email, he did not set up a vacation daemon on his work machine. What happens when an email is received at John's work email address?
 - when an email is received by the mail server attached to John's work email, it is forwarded to the server of the couple's personal email address. The forwarding could have been con-

- figured to retain the envelope sender (plain message forwarding). In this case, an automated vacation would be sent by the personal email mail server back to the sending business. If the forwarding was configured to rewrite the envelope sender as John's work email, then John would receive a vacation auto-reply from the personal work email.
- end up in loop where mail is bouncing back and forth
4. Consider a situation in which a cyberterrorist makes all the DNS servers in the world crash simultaneously. How does this change people's ability to use the Internet?
 - assuming this includes all DNS caches, this would drastically affect the ability of people to use the Internet, as e.g. HTTP requests only use hostnames and rely on DNS resolution to be able to send the request. Services with static IPs which are widely known would be able to continue use
 5. Is it possible for an organizations Web server and mail server to have exactly the same alias for a hostname (for example, foo.com)? How does a DNS server distinguish these two types of mapping, namely, hostname-web server IP and hostname-mail server IP?
 - Yes: the DNS client requests the MX record from the resource record for a given hostname to get the mail server, and recursively searches until and requests until it hits an A record with an IP for the web Server
 6. DNS uses UDP instead of TCP. If a DNS packet is lost, there is no automatic recovery. Does this cause a problem, and if so, how is it solved?
 - DNS will attempt to repeat query, attempt a query to another name server, or or notify application the query was unsuccessful
 7. Suppose you can access the caches in the local DNS servers of your department. Can you propose a way to roughly determine the Web servers (outside your department) that are most popular among the users in your department? Explain.
 - count number of times local DNS servers receive a request for a domain before they are forwarded to the DNS hierarchy (presumably if the entry is cached these won't be forwarded to DNS hierarchy)
 - count length of time between domain being added to the cache: the most popular ones will be close to the TTL, while the less popular ones may have a period longer than the TTL