Network Services

Dr. Manfred Hauswirth
Laboratoire de Systèmes d'Information Répartis
École Polytechnique Fédérale de Lausanne (EPFL)
http://lsirpeople.epfl.ch/hauswirth/

Oct 21, 2002 – Lecture 1

Timetable

- Block, 5 times, 150 min (s.t.!!)

Date	Time	Room	Content
21.10.2002	10:00-12:30	HS 11	1. VO
21.10.2002	14:00-16:00	Radinger-HS	Lab assignment
22.10.2002	16:00-18:30	HS 8	2. VO
23.10.2002	12:00-14:30	FH 8	3. VO
24.10.2002	10:30-13:00	Radinger-HS	4. VO
25.10.2002	10:00-12:30	HS 8	5. VO
18.12.2002			Exam
18.12.2002	24:00		Deadline for Lab

http://www.infosys.tuwien.ac.at/Teaching/Courses/NetworkServices/

Practical part (UE)

- New curriculum only
- 1 example/student
  - Study literature and write a report or
  - Do some programming
- Assignment: Oct 21, 2002, 14:00-16:00, Radinger-HS
  - Enrollment (if you are not yet)
  - More info
  - Assignment of topics
  - **Being there is required!!!**
- Dec 18, 2002, 24:00 – Deadline for lab
  No extensions!

Contents of the lecture (1/2)

- Basic services and protocols
  Telnet, FTP, Email (SMTP, POP, IMAP), Usenet News (NNTP), WWW (HTTP)
- WWW: servers and infrastructure
  planning, configuration, CGI, caches
- Search engines and directory services
  DNS, Glimpse, Harvest, LDAP, robots (if time)
- Standard languages of the web
  HTML, XML, XSL, CSS, DOM
- A little bit of "WWW programming languages"
  Java, JavaScript

Contents of the lecture (2/2)

- Security
  PGP, SSL, public key infrastructures, certificates, firewalls, copyright, data protection, privacy, Java security
- Peer-to-peer systems
  Napster, FastTrack, Gnutella, Freenet, P-Grid
- Push systems / event-based systems
  Foundations, technologies
- E-commerce
  Macropayment vs. micropayment, SET, Millicent
- Overview of mobile code
  Foundations, technologies, security

Exam

- Dec 18, 2002
- Closed book exam
- Content of the lecture
- Enrollment mandatory
  - Web or
  - Argentinierstr. 8, 2nd floor

Up-to-date Information

Network Services homepage

http://www.infosys.tuwien.ac.at/Teaching/Courses/NetworkServices/

Overview of this lecture

- Accessing remote systems: Telnet
- File Transfer: FTP
- Electronic mail: SMTP, POP, IMAP, MIME
- Secure access and transfer: SSH
- Usenet News
- WWW basics

Telnet

- Batch ⇨ interactive working
- "Internet terminal"
  `telnet <host> [<port>]` (default port: 23)
- Catalogs, databases, services, etc. via Telnet ⇨ primitive information systems
  e.g., **`telnet pac.carl.org`**
  Colorado Alliance of Research Libraries (CARL)

Telnet process model

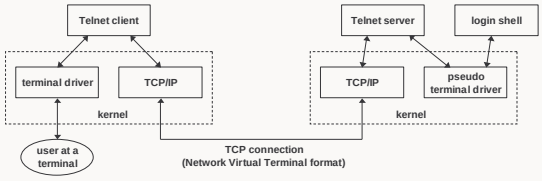


```
telnet pent224.infosys.tuwien.ac.at
Red Hat Linux release 7.1 (Manhattan)
Kernel 2.4.2-2smp on an i686
login: pooh
Password:
tcsh>
```

(hehe - guess it :-)

Telnet's pros and cons

+ Standardized (RFC 854)
+ Debug text-based protocols (HTTP, SMTP, NNTP, POP, ...)
– 1 keystroke = 3 TCP packets (↔)
– Passwords in the clear (like most other remote login tools)
– Data in the clear

Example: Telnet - HTTP

```
tcsh > telnet www.infosys.tuwien.ac.at 80
Trying...
Connected to pent21.infosys.tuwien.ac.at.
Escape character is '^]'.
GET / HTTP/1.0

HTTP/1.1 200 OK
Date: Wed, 01 Sep 2001 18:12:13 GMT
Server: Apache/1.3.14 (Unix) tomcat/1.0 PHP4.0.3pl1 mod_ssl/2.7.1 OpenSSL/0.9.4
Last-Modified: Wed, 28 Aug 2001 10:50:42 GMT
ETag: "ee04d-14bd-3844fd82"
Accept-Ranges: bytes
Content-Length: 5309
Connection: close
Content-Type: text/html

<TITLE>  Distributed Systems Group Home Page </TITLE>
<BODY BGCOLOR=#FFFFFF  TEXT=#000000
  LINK=#0000FF  ALINK=#000000  VLINK=#800080 >
<table border="0" width="100%" cellspacing="0" cellpadding="0">
  <tr>
     <td bgcolor="#000099" VALIGN="CENTER">
```

File Transfer Protocol (FTP)

- Bi-directional file transfer (binary, ASCII)
  `ftp <host>` (default port: 21)
- Interactive access
  - file
  - directory commands
- Authentication via username/password
- Anonymous FTP:
  login as **`anonymous/<email address>`**

FTP process model

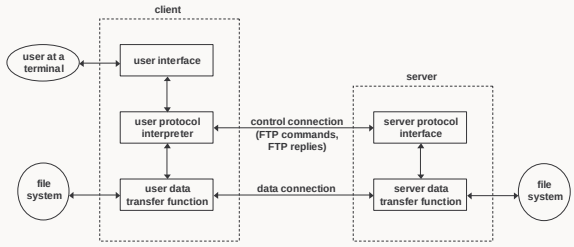

FTP sample session

```
tcsh > ftp swd.infosys.tuwien.ac.at
Connected to swd.infosys.tuwien.ac.at.
220 swd.infosys.tuwien.ac.at FTP server ... Tue Feb 14 14:38:22 ... ready.
Name (swd.infosys.tuwien.ac.at:pooh): anonymous
331 Guest login ok, send your complete e-mail address as password.
Password:
230 Guest login ok, access restrictions apply.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> cd pub/redhat71/images
250 CWD command successful.
ftp> bin
200 Type set to I.
ftp> get boot.img
local: boot.img remote: boot.img
200 PORT command successful.
150 Opening BINARY mode data connection for boot.img (1474560 bytes).
226 Transfer complete.
1474560 bytes received in 57.9 secs (25 Kbytes/sec)
ftp> quit
tcsh >
```


FTP commands

- User commands
  - get, put, mget, mput
  - cd, dir, lcd
  - prompt, ascii, binary
- Protocol commands
  - USER, PASS, QUIT
  - RETR, STOR
  - CWD, LIST

FTP - Telnet

```
tcsh > telnet swd.infosys.tuwien.ac.at 21
Trying 128.131.172.90...
Connected to swd.infosys.tuwien.ac.at.
Escape character is '^]'.
220 swd.infosys.tuwien.ac.at FTP server (Version wu-2.4.2-academ[BETA-18-
    VR14-SKEY](1) Tue Feb 14 14:38:22 PST 2001) ready.
USER anonymous
331 Guest login ok, send your complete e-mail address as password.
PASS pooh@infosys.tuwien.ac.at
230 Guest login ok, access restrictions apply.
CWD pub/redhat71/images
250 CWD command successful.
LIST
425 Can't build data connection: Connection refused.
QUIT
221 Goodbye.
Connection closed by foreign host.
```


FTP replies

- 3-digit code (ASCII) + optional message
- Example replies:
  125 Data connection already open; transfer starting.
  200 Command OK.
  331 Username OK, password required.
  425 Can't open data connection.
  500 Syntax error (unrecognized command).
- QUIT ⇨ 221 Goodbye

FTP's pros and cons

+ Standardized (RFC 959)
+ Anonymous FTP

– Passwords in the clear
– Data in the clear

Electronic Mail

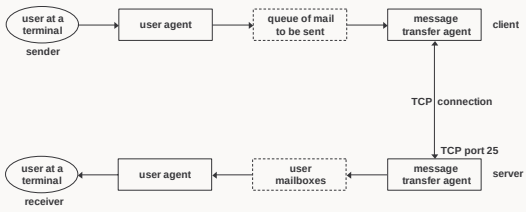

Protocols for email

- Simple Mail Transfer Protocol (SMTP)
  - mail exchange between MTAs
  - default port: 25
- Post Office Protocol (POP)
  - retrieve mail from a mail drop
  - default port: 110 (POP3)
- Internet Message Access Protocol (IMAP)
  - access and manage remote message stores
  - default port: 143

SMTP sample session

```
tcsh > mail -v ailec@infosys.tuwien.ac.at
Subject: test
just a test
.
ailec@infosys.tuwien.ac.at... Connecting to mail.infosys... via nullclient...
220 infosys.tuwien.ac.at ESMTP Sendmail 8.9.3; Thu, 2 Sep 2001 12:21:13 ...
HELO pent224.infosys.tuwien.ac.at
250 infosys.tuwien.ac.at Hello pent224.infosys... [...], pleased to meet you
MAIL From:<pooh@infosys.tuwien.ac.at>
250 <pooh@infosys.tuwien.ac.at>... Sender ok
RCPT To:<ailec@infosys.tuwien.ac.at>
250 <ailec@infosys.tuwien.ac.at>... Recipient ok
DATA
354 Enter mail, end with "." on a line by itself
just a test
.
250 MAA12507 Message accepted for delivery
QUIT
221 infosys.tuwien.ac.at closing connection
```


The corresponding email

Email format specified in RFC 821 and RFC 822

```
Return-Path: pooh@infosys.tuwien.ac.at
Delivery-Date: Thu Sep 2 12:24 MET 2001
Received: from pent224.infosys.tuwien.ac.at (pent224.infosys.tuwien.ac.at
   [128.131.172.124])
   by infosys.tuwien.ac.at (8.9.1b+Sun/8.9.3) with ESMTP id MAA12517
   for <ailec@infosys.tuwien.ac.at>; Thu, 2 Sep 2001 12:24:10 +0100 (MET)
Received: (from pooh@localhost)
   by pent224.infosys.tuwien.ac.at (8.8.7/8.8.5) id MAA31796
   for ailec@infosys.tuwien.ac.at; Thu, 2 Sep 2001 12:25:40 +0100
Message-Id: <200112021125.MAA31796@pent224.infosys.tuwien.ac.at>
From: Manfred Hauswirth <pooh@infosys.tuwien.ac.at>
Date: Thu, 2 Sep 2001 12:25:40 +0100
To: ailec@infosys.tuwien.ac.at
Subject: test

just a test
```


SMTP commands

- Minimal SMTP command set (RFC 821)
  - HELO client identification
  - MAIL originator
  - RCPT email address of message recipient
  - DATA message contents
  - QUIT terminate mail exchange
  - RSET abort current mail transaction
  - VRFY verify recipient address
  - NOOP do nothing

SMTP replies

- 3-digit code (ASCII) + optional message
- Example replies:
  - 250 <address>... Sender ok
  - 354 Enter mail, end with "." on a line by itself
  - 421 <domain> Service not available, closing transmission channel
  - 500 Syntax error (unrecognized command).
- QUIT ⇨ 221 <domain> closing connection

Extended SMTP

Extended SMTP (ESMTP): RFC 1425

```
220 infosys.tuwien.ac.at ESMTP Sendmail 8.9.1b+Sun/8.9.3; Thu, 2 Sep 2001
   17:40:26 +0100 (MET)
EHLO w4.infosys.tuwien.ac.at
250-infosys.tuwien.ac.at Hello w4.infosys.tuwien.ac.at [128.131.172.44],
   pleased to meet you
250-EXPN          expand mailing list (do not send mail)
250-VERB
250-8BITMIME      8 bit characters supported in mail body
250-SIZE          do you accept a mail of size <x> ?      (RFC 1427 !!)
250-DSN
250-ONEX
250-ETRN
250-XUSR
250 HELP
```


Mail eXchange (MX) records

- Special resource record in the DNS
- Define mail relay for domains (hosts)

```
tcsh > nslookup
> set qt=mx
> sun.com
Server:  w7.infosys.tuwien.ac.at
Address:  128.131.172.47
sun.com preference = 40, mail exchanger = mondzo.sun.com
sun.com preference = 5, mail exchanger = venus.Sun.COM
sun.com preference = 5, mail exchanger = lukla.Sun.COM
sun.com preference = 15, mail exchanger = mercury.Sun.COM
sun.com preference = 40, mail exchanger = mars.sun.com
mondzo.sun.com  internet address = 192.18.100.1
venus.Sun.COM   internet address = 192.9.25.5
.......
>
```


Anatomy of an email

```
Return-path: eXXXXXXX@stud4.tuwien.ac.at
Delivery-date: Mon Aug 10 10:55 MET 2001
Received: from stud4.tuwien.ac.at (stud4.tuwien.ac.at [193.170.75.21])
   by infosys.tuwien.ac.at (8.9.1b+Sun/8.9.3) with ESMTP id KAA15892
   for <pooh@infosys.tuwien.ac.at>; Mon, 8 Aug 2001 10:55:47 +0100 (MET)
Received: from stud4.tuwien.ac.at (eXXXXXXX@stud4.tuwien.ac.at [193.170.75.21])
   by stud4.tuwien.ac.at (8.9.3/8.9.3) with ESMTP id KAA14576
   for <pooh@infosys.tuwien.ac.at>; Mon, 8 Aug 2001 10:56:58 +0100 (MET)
Message-id: <Pine.HPX.4.10.9911081041390.6983-100000@stud4.tuwien.ac.at>
In-reply-to: Your message of "Thu, 7 Aug 2001 15:36:32 +0100."
             <199911111446.PAA10875@w4.infosys.tuwien.ac.at>
Date: Wed, 8 Aug 2001 10:56:58 +0100 (MET)
Reply-to: eXXXXXXX@student.tuwien.ac.at
Precedence: normal
Sender: eXXXXXXX@stud4.tuwien.ac.at
From: John Doe <eXXXXXXX@student.tuwien.ac.at>
To: pooh@infosys.tuwien.ac.at
Subject: Re: Proposal for Master's Thesis

Hi Manfred!
.........
```


Post Office Protocol

- Retrieve mail from a mail drop
- Default port: 110 (POP3)

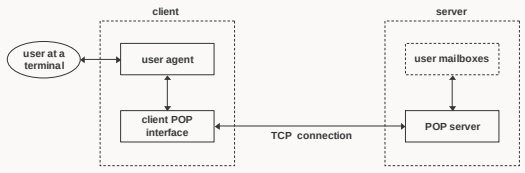

POP sample session

```
+OK POP3 pop3.infosys.tuwien.ac.at v7.59 server ready
USER pooh
+OK User name accepted, password please
PASS ratzfatz
+OK Mailbox open, 2 messages
LIST
+OK Mailbox scan listing follows
1 1485
2 2281
.
RETR 1
+OK 1485 octets
Received: from w4.infosys.tuwien.ac.at (w4.infosys.tuwien.ac.at ....
     the mail
.
DELE 1
+OK Message deleted
DELE 3
-ERR No such message
QUIT
+OK Sayonara
Connection closed by foreign host.
```

POP commands and replies (RFC 1939)

- Main commands
  - USER specify user name (APOP name digest)
  - PASS user's password (in the clear !!!)
  - QUIT end session
  - STAT number of messages, size of mailbox
  - LIST list mailbox/message, size & number of message
  - RETR retrieve a message
  - DELE mark message for deletion
  - NOOP do nothing
  - RSET unmark message for deletion
- Replies
  - +OK [<text>]
  - -ERR [<text>]

Internet Message Access Protocol

- POP's shortcomings
  - password in the clear (APOP not implemented)
  - transfer mail to local system ⇨ network access? Mail scattered on several computers.
  - leave mail on POP server ⇨ unstructured
- IMAP addresses these shortcomings and supports access and management of email on a mail server (RFC 2060 + Draft)

IMAP goals

- Keep mail on the mail server: mail needs not be transferred back and forth
- Manipulation of remote mailboxes as if they were local
- Access & management from > 1 computer
  - full access to mail while travelling
  - consistent access from home, work, etc.
- Access modes: on-/offline, disconnected
- Concurrent access to shared mailboxes

IMAP access modes

- Offline: similar to POP
- Online: access, manage and manipulate mail on a server (mail is left on the server)
- Disconnected (mail is left on the server)
  - make a cache copy of selected messages and disconnect
  - manipulate cache copy
  - reconnect to the server and resynchronize

IMAP > POP

- Access and management of multiple mailboxes
- Concurrent access and update to shared mailboxes
- Standard and user-defined status flags (seen, answered, draft, ...)
- Manipulate remote folders != INBOX
- List/create/delete/rename remote folders
- Support for folder hierarchies
- Append messages to remote folder
- Determine message structure without downloading
- Selective fetching of individual MIME body parts
- Server-based searching and selection to minimize data transfer
- Negotiated extensions => extend capabilities

IMAP problems

- Security holes (most of them are fixed)
- Stability (pretty good now)
- Online mode: no IMAP server => no access to "old" mail
- Complex protocol => hard to implement
- Mail clients still support POP rather than IMAP
- Password in the clear

Multipurpose Internet Mail Extensions (MIME)

- Mail body (RFC 822)
  - 7 bit ASCII
  - max line length: 1000 (SMTP DATA command)
- MIME (RFC 2045, RFC 2046, RFC 2047, RFC 2048, RFC 2049)
  - overcome RFC 822 limitations
  - requires no change to MTAs
  - structured mail content

MIME header fields (1/3)

- MIME-Version
  - version of the Internet message body format standard in use
  - **`MIME-version: 1.0`**
- Content-Type
  - describe the data contained in the body => user agent can pick an appropriate agent to present the data to the user
  - **`Content-Type: text/plain`**

MIME header fields (2/3)

- Content-Transfer-Encoding
  - describe encoding of binary data (non-7-bit) into a 7 bit short line format (SMTP, RFC 821)
  - **`Content-Transfer-Encoding: base64`**
- Content-ID
  - one body may want to reference another => unique id required
  - **`Content-ID: <id42@w6infosys.tuwien.ac.at>`**

MIME header fields (3/3)

- Content-Description
  - associate some descriptive information with a given message body
  - **`Content-Description: Mail message body`**

MIME example (screenshot)

MIME example (raw content)

```
Subject: test
Content-Type: multipart/mixed; boundary="border_1"

This is a multi-part message in MIME format.
--border_1
Content-Type: multipart/alternative; boundary="border_2"

--border_2
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
Content-Description: Mail message body

Just a nice test mail
--border_2
Content-Type: text/html; charset=us-ascii
Content-Transfer-Encoding: 7bit

<HTML>
Just a nice test mail
</HTML>
--border_2--

--border_1
Content-Type: image/gif; name="creperie.gif"
Content-Transfer-Encoding: base64

R0lGODlhIQHRAPcAAAAAIAAAACAAICAAAAAgIAAgACAgICAgAQEBPz8/AAAAAAAAAAAAAA
....
--border_1--
```

Some MIME content types

Content Type	Subtype
text	plain html xml
multipart	mixed alternative
application	octet-stream postscript
image	gif jpeg
audio	basic midi mpeg

Protocol security problems

- Telnet, FTP, POP, etc. send passwords in the clear
  - login password used
  - password can be intercepted and misused
- Telnet, FTP, POP, etc. send data in the clear (messages)
  - data can be intercepted and misused

Secure Shell (SSH)

- Automatic authentication of users
  - passwords are encrypted => no password sniffing possible
- Authentication of both connection ends
  - server and client are authenticated => identity of communication partners
- Encryption and compression of data
  - => secure and fast transfer of data
  - tunneling and encryption of arbitrary connections (protocols)
- Problem: How to obtain certificates?

Additional SSH benefits

- Virtual Private Networks (VPNs) via SSH
- No IP spoofing / IP source spoofing
  (a host sends out packets which pretend to come from a trusted host)
- No DNS spoofing
  (an attacker forges DNS records)
- No data manipulation on intermediate hosts
- Protection of X11 authentication data and against spoofed X11 connections

SSH tunneling of connections

```
myhost> ssh -L 1234:pophost:110 pophost sleep 100
```

```
myhost> telnet myhost 1234
+OK POP3 pophost v7.59 server ready
USER pooh
+OK User name accepted, password please
PASS ratzfatz
+OK Mailbox open, 2 messages
....
```

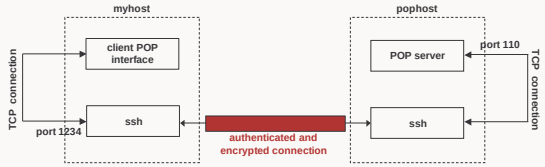

FTP via SSH

- FTP server must accept PORT command
- FTP client must use PORT command
- does not work with wu-ftpd:

```
myhost> ssh -L 1234:ftphost:21 ftphost sleep 100
```

```
myhost> ftp myhost 1234
Connected to localhost.
220 localhost FTP server (SunOS 5.7) ready.
Name (localhost:pooh):
331 Password required for pooh.
Password:
230 User pooh logged in.
ftp> dir
200 PORT command successful.
425 Can't build data connection: Connection refused.
```

SSH commands

- ssh [-l login] host [command ...]
  - remote login, remote command execution
  - replacement for rsh
- slogin host
  - remote login
  - replacement for rlogin, telnet
- scp [[user@host1:]file1 ... [[user@host2:]file2
  - remote copy
  - replacement for rcp, ftp

What is (Usenet) News ?

- "A global distributed blackboard on top of other networks."
- "A World-wide discussion forum which is divided into hierarchical newsgroups dedicated to defined topics."

News basics

- Approximately 60000 newsgroups (~ topics)
- Newsgroups are hierarchically structured (dot notation)
  e.g., at.tuwien.student, comp.lang.java
- Users can subscribe to a set of newsgroups
- Users can write new articles (or postings) and post them to a (set of) newsgroup(s) ⇨ all subscribers see the postings
- Users can reply to existing postings ⇨ thread (discussion)

Newsreader screenshot

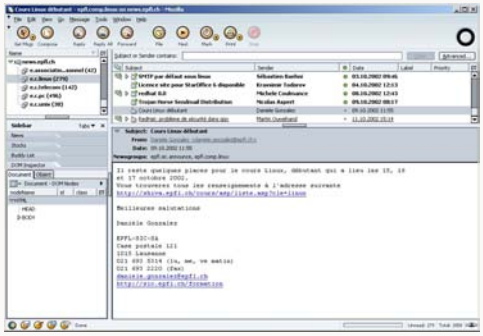

Newsgroup hierarchies

Name	Topic
alt	misc. topics
biz	business related topics
comp	computers
rec	recreation
sci	science
soc	society
talk	general talk

comp.lang	programming languages
comp.lang.java	the Java language
comp.lang.security java	security issues of Java

Types of Newsgroups

- Reading and posting allowed
  - everybody can read and post articles
- Moderated
  - everybody can read articles
  - articles to be posted are sent to the moderator who decides whether they are being posted
- Read-only
  - everybody can read articles
  - posting requires special authorization

Articles

- Articles can be posted to one ore more newsgroups (crossposting)
- If an article is crossposted, it should name a "Follow-Up" newsgroup to channel responses
- "Protocol Data Units"

Anatomy of an article

```
Path: news.tuwien.ac.at!news.univie.ac.at
From: Manfred Hauswirth <M.Hauswirth@infosys.tuwien.ac.at>
Newsgroups: at.tuwien.student, at.tuwien.general
Subject: LVA-Ankuendigung: Network Services VO 1.0 (184.023)
Followup-To: at.tuwien.student
Date: Fri, 19 Nov 1999 11:24:38 +0100
Organization: Technical University of Vienna, Distributed Systems Group
Lines: 68
Message-ID: <38352565.AE1A78D8@infosys.tuwien.ac.at>
NNTP-Posting-Host: darkstar.infosys.tuwien.ac.at

Network Services VO 1.0 (184.023)

WS 1999/2000 - Manfred Hauswirth

Lehrziele

    Ziel dieser Vorlesung ist es, ...
```

unique id (do not re-use) identify duplicate article offers

News distribution

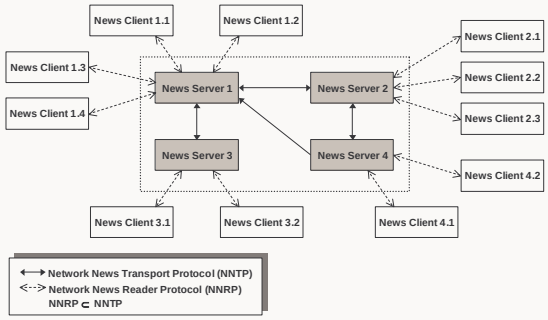

News Servers

- Distribute articles and admin info (new newsgroups) via NNTP
- Offer and distribute a certain subset of all available newsgroups
- Clients can only subscribe to offered newsgroups (via NNRP ⊂ NNTP)
- Index and expire articles (disk space !)
- Typical newsfeed: 2.5GB/day (400kBit/s)
- Standard software: INN (InterNetNews)

Network News Transport Protocol

- Standardized in RFC 977
- Stream-based (TCP)
- SMTP-like commands and responses
- Example commands
  - {ARTICLE, HEAD, BODY} [msgID|#]
  - GROUP <newsgroup>, LIST, LISTGROUP
  - POST, NEWNEWS

NNTP Respones

- 3 digit code (ASCII) + optional message
- Last line of multi-line responses: "."

Reply	Description
1yz	Informative message
2yz	Command OK
3yz	Cmd OK so far; send rest of cmd
4yz	Cmd correct but failed
5yz	Error

Reply	Description
x0z	Connection, setup, misc. message
x1z	Newsgroup selection
x2z	Article Selection
x3z	Distribution function
x4z	Posting
x8z	Nonstandard extensions
x9z	Debugging output

NNTP sample session

```
200 news.tuwien.ac.at INN NNRP server INN 1.7.2 ... ready (posting ok).
ARTICLE 13061
412 Not in a newsgroup
GROUP at.tuwien.xxx
411 No such group at.tuwien.xxx
WRONGCOMMAND 123
500 What?
GROUP at.tuwien.student
211 2141 10786 13260 at.tuwien.student
XOVER 13061
224 data follows
13061   LVA-Ankuendigung: Network Services ... 19 Nov 1999 11:24:38 +0100 ...
   <38352565.AE1A78D8@infosys.tuwien.ac.at>  2840  68 Xref: news.tuwien.ac.at
   at.tuwien.student:13061
ARTICLE 13061
220 13061 <38352565.AE1A78D8@infosys.tuwien.ac.at> article
Path: news.tuwien.ac.at!tuwien-news
From: Manfred Hauswirth <M.Hauswirth@infosys.tuwien.ac.at>
   ...
.
QUIT
205 .
```


News Problems

- Delivery is not guaranteed (best effort)
- Reply may be received before article
- Scalability
  - every article is copied to all other News servers
  - ⇨ fast response :-) few articles are read :-(
  - ⇨ high bandwidth (400kBit/s) and disk space consumption (2.5GB/day, 1 file/article)
  - ⇨ backlog, special HW, disks for News server
  - ⇨ cache servers

like email, mailing lists !!

A brief World Wide Web History

1945	V. Bush: "As we may think" (memex, linked info)
1980	T. Berners-Lee: Enquire system
1989	T. Berners-Lee: Proposal for a global hypertext system (CERN)
1991	T. Berners-Lee: 1st browser/editor (CERN)
1993	M. Andreessen: Mosaic, NCSA server (NCSA)
1994	Netscape is founded
1996	Microsoft "invents" the Internet
1997	The browser war starts
1996-?	Myriads of languages, protocols, standards
Now	WWW + cell phones + PDAs + ...

WWW Conceptually

- The WWW is a synchronous, distributed, multiple client, multiple server hyper-media system.
  - HyperText Transfer Protocol (HTTP)
  - HyperText Markup Language (HTML)

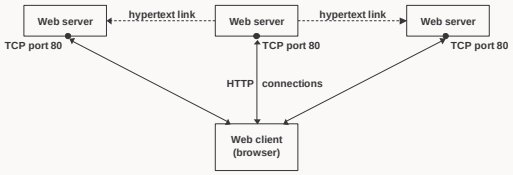

Uniform Resource XYZ

- URI: Naming scheme and syntax
- URL: URIs for objects accessed with existing protocols
- URN: persistant object names (resolution protocols)
- Hyperlinks on the WWW are represented as URLs

Uniform Resource Locator

- <scheme>:<scheme-specific-part>
- Schemes
  - ftp, http, mailto, news, telnet, file, ...
- Common Internet Scheme Syntax
  //[<user>[:<password>]@]<host>[:<port>]/[<path>[#<fragment>]]

URL Examples

http://www.dslab.tuwien.ac.at/Lab3/homepage.html#Hints
http://wp.tuwien.ac.at:8888/Ro%3dTechnische%20Universitaet%20Wien%2c%20c%3dAT
ftp://gd.tuwien.ac.at/
ftp://anonymous@gd.tuwien.ac.at/gnu/
ftp://anonymous:@gd.tuwien.ac.at/gnu/
ftp://hugo:magumba@gd.tuwien.ac.at/gnu/
mailto:manfred.hauswirth@epfl.ch
news://news.tuwien.ac.at/at.tuwien.general
news://news.tuwien.ac.at/at.tuwien.general
telnet://hugo@brio07.infosys.tuwien.ac.at
file:///home/users/hugo/file1.txt

HTTP Interaction Pattern

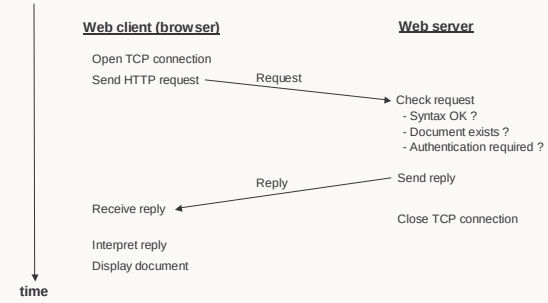

Ahead

- World-wide Web
  - HTTP
  - Servers
- Advanced WWW Concepts
  - CGI
  - Cookies
  - SSI
  - Servlets
  - Caching