

# The Application Layer: HTTP, SMTP

Lecture 6

<http://www.cs.rutgers.edu/~sn624/352-S22>

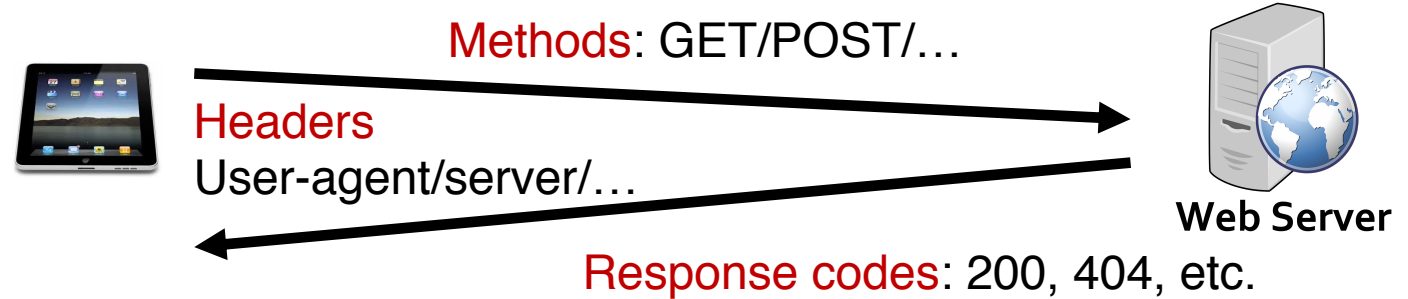
Srinivas Narayana

# Quick recap of concepts




## HyperText Transfer Protocol (HTTP)


HTTP is a client/server application



## Persistence



## Cookies



**Caching**  
e.g., proxy  
server

# Web caches

Web caches: Machines that remember web responses for a network

## Why cache web responses?

- Reduce response time for client requests
- Reduce traffic on an institution's access link

Last lecture: Caches implemented in the form of a **proxy server**

# Content Distribution Networks (CDNs)

## A global network of web caches

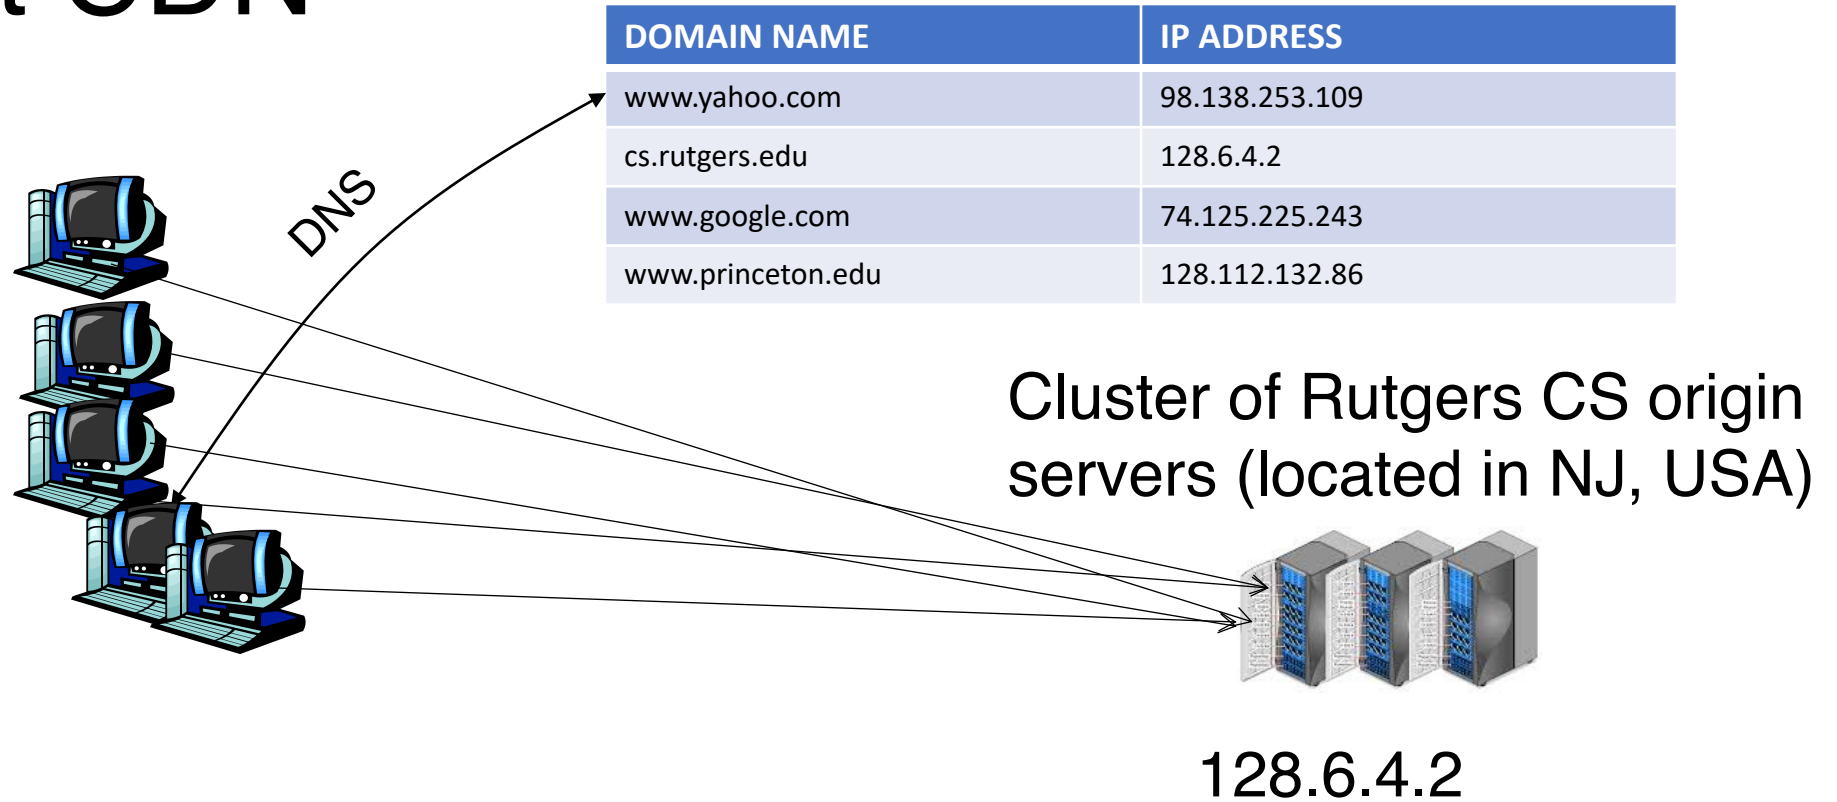
- Provisioned by ISPs and network operators
- Or content providers, like Netflix, Google, etc.

Uses (overlaps with uses of web caching in general)

- Reduce traffic on a network's Internet connection, e.g., Rutgers
- Improve response time for users: CDN nodes are closer to most users than origin servers
- Reduce bandwidth requirements on content provider
- Reduce \$\$ to maintain origin servers

# Without CDN

Clients  
distributed  
all over the  
world



- Problems:
- Huge bandwidth requirements for Rutgers
- Large propagation delays to reach users

# Where the CDN comes in

- Distribute content of the origin server over geographically distributed **CDN servers**
- But how will users get to these CDN servers?
- **Use DNS!**
  - DNS provides an additional layer of indirection
  - Instead of returning IP address, return another DNS server (NS record)
  - The second DNS server (run by the CDN) returns IP address to client
- The CDN runs its own DNS servers (**CDN name servers**)
  - Custom logic to send users to the “closest” CDN web server

# With CDN

NS record delegates the choice of IP address to the CDN name server.

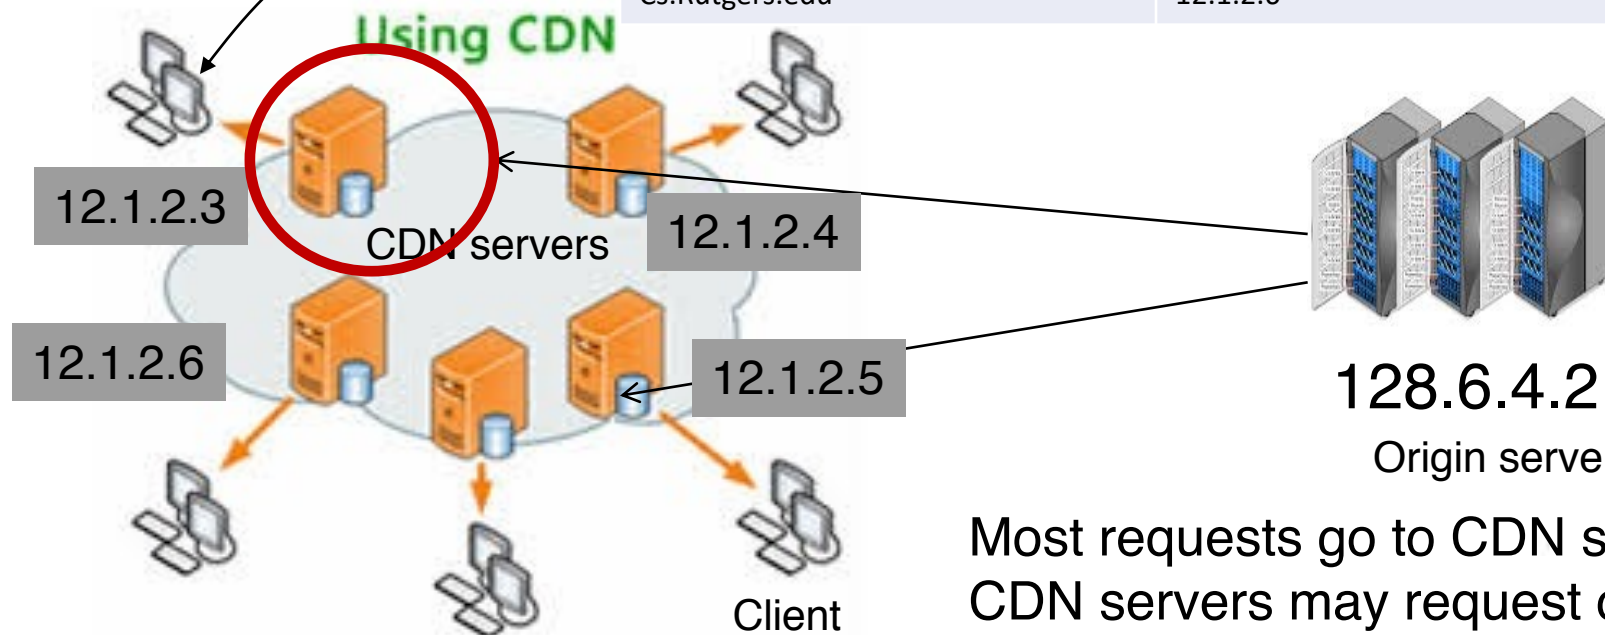
DOMAIN NAME	IP ADDRESS
www.yahoo.com	98.138.253.109
cs.rutgers.edu	124.8.9.8 (NS record pointing to CDN name server)
www.google.com	74.125.225.243

## CDN Name Server (124.8.9.8)

DOMAIN NAME	IP ADDRESS
Cs.Rutgers.edu	12.1.2.3
Cs.Rutgers.edu	12.1.2.4
Cs.Rutgers.edu	12.1.2.5
Cs.Rutgers.edu	12.1.2.6

Custom logic to map ONE domain name to one of many IP addresses!

Popular CDNs:  
CloudFlare  
Akamai  
Level3  
...



Most requests go to CDN servers (caches).  
CDN servers may request object from origin  
Few client requests go directly to origin server

# Summary of HTTP



- Request/response protocol
- ASCII-based human-readable message structures
- Improve performance using connection persistence, caching, and CDN
- Enhanced stateful functionality using cookies
- Simple, highly-customizable protocol
  - Just add headers
- Protocol that forms of the basis of the web we enjoy today!





# Simple Mail Transfer Protocol

# We're all familiar with email. How does it work?




# Electronic Mail

## Three major components:

### 1. User agents

- a.k.a. “mail reader”
- e.g., Applemail, Outlook
- Web-based user agents (ex: gmail)




# Electronic Mail: Mail servers

## 2. Mail Servers

- Mailbox contains incoming messages for user
- Message queue of outgoing (to be sent) mail messages
- Sender's mail server makes connection to Receiver's mail server
  - IP address, port 25

## 3. SMTP protocol: client/server protocol

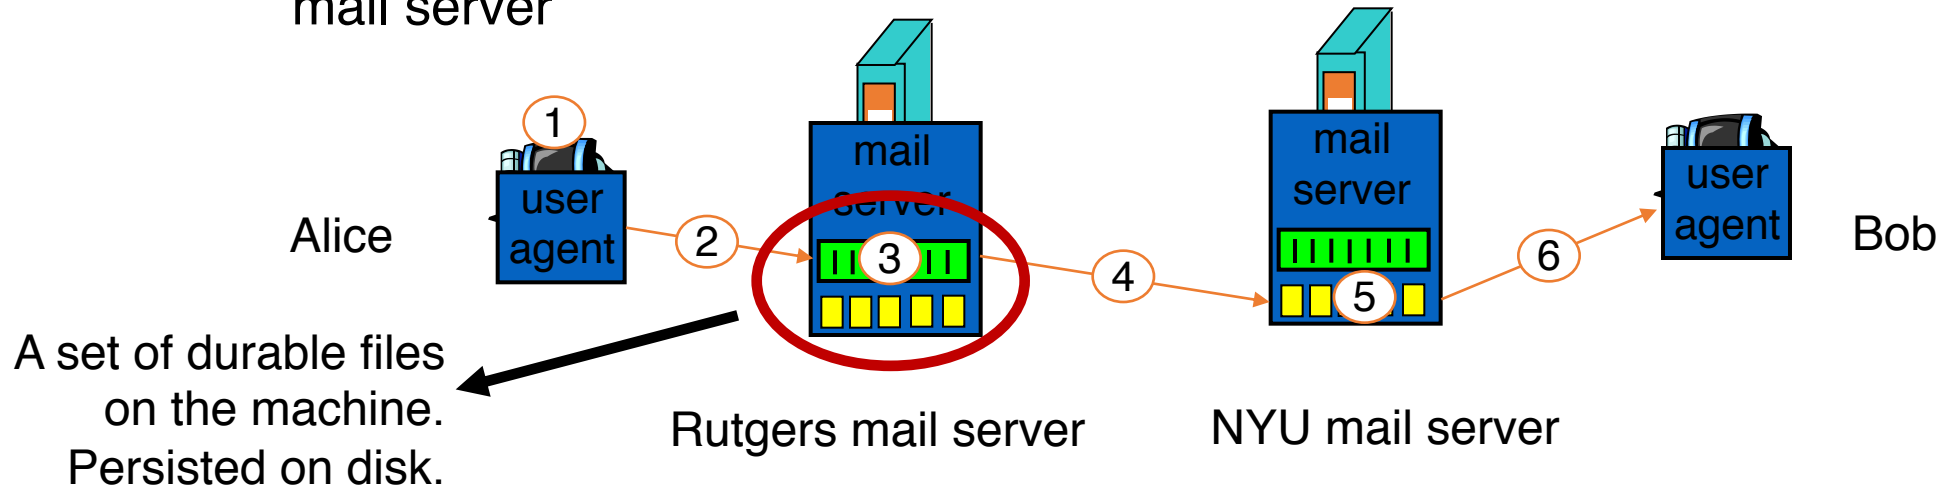
- Used to **send** messages
- Client: sending user agent or sending mail server
- server: receiving mail server



# Scenario: Alice sends message to Bob

- 1) Alice (alice@rutgers.edu) uses UA to compose message to bob@nyu.edu
- 2) Alice's UA sends message to her mail server; message placed in outgoing message queue
- 3) Client side of SMTP opens TCP connection with Bob's mail server

- 4) SMTP client sends Alice's message over the TCP connection
- 5) Bob's mail server places the message in Bob's incoming mailbox
- 6) Sometime later, Bob invokes his user agent to read message



# Observations on these exchanges

- Mail servers are the “infrastructure” for email functionality
  - Receiving the email on behalf of Bob, should Bob’s machine be turned off
  - Retrying the delivery of the email to Bob on behalf of Alice, should Bob’s mail server be unavailable in the first attempt
- The same machine can act as client or server based on context
  - Rutgers’s mail server is the server when Alice sends the mail
  - It is the client when it sends mail to Bob’s mail server
- SMTP is push-based: info is pushed from client to server
  - Contrast to HTTP or DNS where info is pulled from the server

# Sample SMTP interaction

- telnet <mail-server> 25
- HELO <sender-domain>
- MAIL FROM: <name@<sender-domain>- RCPT TO: <user>@<mail-server-domain>
- DATA
- Put data in, then [enter].[enter] Don't forget the "."
- You can add mail headers (later) to make your email look good



# MAIL command response codes

**Table 23.2** *Responses*

<i>Code</i>	<i>Description</i>
<b>Positive Completion Reply</b>	
<b>211</b>	System status or help reply
<b>214</b>	Help message
<b>220</b>	Service ready
<b>221</b>	Service closing transmission channel
<b>250</b>	Request command completed
<b>251</b>	User not local; the message will be forwarded
<b>Positive Intermediate Reply</b>	
<b>354</b>	Start mail input
<b>Transient Negative Completion Reply</b>	
<b>421</b>	Service not available
<b>450</b>	Mailbox not available
<b>451</b>	Command aborted: local error
<b>452</b>	Command aborted; insufficient storage
<b>Permanent Negative Completion Reply</b>	
<b>500</b>	Syntax error; unrecognized command
<b>501</b>	Syntax error in parameters or arguments
<b>502</b>	Command not implemented
<b>503</b>	Bad sequence of commands
<b>504</b>	Command temporarily not implemented
<b>550</b>	Command is not executed; mailbox unavailable
<b>551</b>	User not local
<b>552</b>	Requested action aborted; exceeded storage location
<b>553</b>	Requested action not taken; mailbox name not allowed
<b>554</b>	Transaction failed

220: Service ready

250: Request command complete

354: Start mail input

421: Service not available

500: Unrecognized command

# Mail message format (stored on server)

SMTP text message exchange  
standardized in RFC 822

- **Header lines**, e.g.,


- To:
- From:
- Subject:

*These are different from SMTP  
commands!*

(these would still be under  
“DATA”)

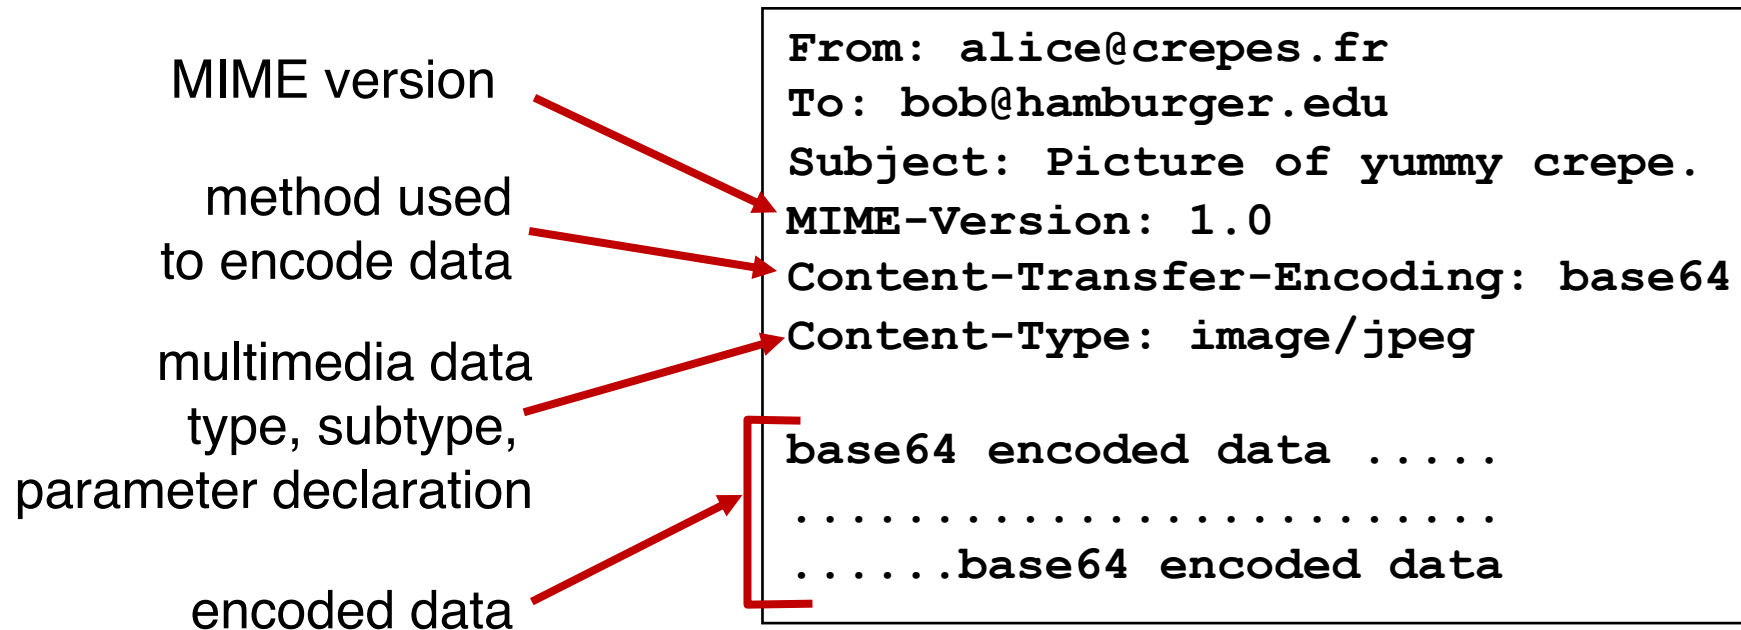
- **body**

- the “message”.
- ASCII characters only



# Message format: multimedia extensions


- MIME: multimedia mail extension, RFC 2045, 2056
- additional headers in DATA header declare MIME content type
- A message can have many parts



Vote for New Brunswick's Favorite Restaurants

Inbox



 @NBCCityCenter citymarket@newbrunswick.com via newbrunswick.ccsend.com to me

1:17 AM (8 hours ago)

- Reply
- Forward
- Filter messages like this
- Print
- Add The Morning Paper to Contacts list
- Delete this message
- Block "The Morning Paper"
- Report spam
- Report phishing
- Show original
- Translate message
- Download message
- Mark as unread



**LIFESTYLE & EVENTS**

Upcoming Fun in New Brunswick



Message-ID: <1138527284182.1100462763056.2128912358.0.1561738JL.2002@scheduler.constantcontact.com>

Date: Wed, 2 Feb 2022 17:41:43 -0500 (EST)

From: "@NBCityCenter" <citymarket@newbrunswick.com>

Reply-To: citymarket@newbrunswick.com

Sender: "@NBCityCenter" <pamela@newbrunswick.ccsend.com>

To: [REDACTED]

Subject: Vote for New Brunswick's Favorite Restaurants

MIME-Version: 1.0

Content-Type: multipart/alternative; boundary="-----=\_Part\_550097535\_1238800096.1643841703083"

List-Unsubscribe: <[https://visitor.constantcontact.com/do?](https://visitor.constantcontact.com/do?p=un&m=001HGjyDxPJyPKvJg5rDoRslg%3D%3D&se=001BAprMYzSEJkLWw3FBEdL0g%3D%3D&t=001EkZLEx15CcE%3D&llr=qblcl7aab)

p=un&m=001HGjyDxPJyPKvJg5rDoRslg%3D%3D&se=001BAprMYzSEJkLWw3FBEdL0g%3D%3D&t=001EkZLEx15CcE%3D&llr=qblcl7aab>

List-Unsubscribe-Post: List-Unsubscribe=One-Click

X-Campaign-Activity-ID: 7af9b840-b29e-4db4-8870-83aaf3a5e6a0

X-250ok-CID: 7af9b840-b29e-4db4-8870-83aaf3a5e6a0

X-Channel-ID: f1703300-eb85-11e8-8729-d4ae5275b3f6

X-Mailer: Roving Constant Contact 2012 (<http://www.constantcontact.com>)

X-Return-Path-Hint: Aevm4QLKeTbSIcIOq86XmoA==\_1100462763056\_8XAzAOuFEeiHKdSuUnWz9g==@in.constantcontact.com

X-Roving-Campaignid: 1138527284182

X-Roving-Id: 1100462763056.2128912358

X-Feedback-ID: f1703300-eb85-11e8-8729-d4ae5275b3f6:7af9b840-b29e-4db4-8870-83aaf3a5e6a0:1100462763056:CTCT

X-CTCT-ID: f15ea6d0-eb85-11e8-8729-d4ae5275b3f6

-----=\_Part\_550097535\_1238800096.1643841703083

Content-Type: text/plain; charset=utf-8

Content-Transfer-Encoding: quoted-printable

Dine, Shop, Play & Stay... It All Happens Here! VALENTINE'S WEEKEND =EF=BB=  
=BFLive music & theater entertainment, comedy shows, a variety of activitie=  
s, chocolate shops, and high-end restaurants make New Brunswick an excellen=  
t choice for couples on Valentine's Day! It's easy to see why this city is=  
romantic with so many fun date night options, luxurious spas, and deliciou=  
s dining selections. Book Your Reservations Now Jersey's Choice Restaurant =  
Poll Presented by New Brunswick's Performing Arts Center, it's your time to=  
show your support for your favorite spots in New Brunswick as NJ Monthly h=  
osts their 39th Annual Restaurants of NJ Poll! New Brunswick is represented=  
strongly with many of your favorites including, Roosterspin, Stage Left St=

Constant Contact Data Notice Sent by citymarket@newbrunswick.com powered by Try email marketing for free today!

-----=\_Part\_550097535\_1238800096.1643841703083

Content-Type: text/html; charset=utf-8

Content-Transfer-Encoding: quoted-printable

<!DOCTYPE html>

<html>

<head>

<meta content="text/html; charset=utf-8" http-equiv="Content-Type">

<meta content="width=device-width, initial-scale=1.0" name="viewport" t">

<meta charset="UTF-8">

<meta name="viewport" content="width=device-width, initial-scale=1, maximum-scale=1">

<!--[if gte mso 9]>

<style id="ol-styles">

/\* OUTLOOK-SPECIFIC STYLES \*/

li {

text-indent: -1em;

padding: 0;

margin: 0;

line-height: 1.2;

}

ul, ol {

padding: 0;

margin: 0 0 0 40px;

}

p {



```
    style="color: #333333; text-decoration: none; font-family: cursive;">free today!</a>
```

```
    </div>
```

```
    =20
```

```
    =20
```

```
    </td>
```

```
  </tr>
```

```
</table>
```

```
</td>
```

```
</tr>
```

```
</table>
```

```
</div>
```

```
</td>
```

```
</tr>
```

```
</table></body>
```

```
</html>
```

```
-----=_Part_550097535_1238800096.1643841703083--
```

i:Exit -:PrevPg <Space>:NextPg v:View Attachm. d:Del r:Reply j:Next ?:Help

Date: Wed, 2 Feb 2022 17:41:43 -0500 (EST)

From: "@NBCityCenter" <citymarket@newbrunswick.com>

To:

Subject: Vote for New Brunswick's Favorite Restaurants

X-Mailer: Roving Constant Contact 2012 (<http://www.constantcontact.com>)

[-- Autoview using lynx -assume\_charset='utf-8' -display\_charset=utf-8 -collapse\_br\_tags -dump  
+ '/var/folders/6r/g2lff0zd62q8p\_4sm8h14mqh0000gp/T//mutt.html' --]

Dine, Shop, Play & Stay... It All Happens Here!

[1]Valentine's Day in New Brunswick - Weekend Specials & Events  
VALENTINE'S WEEKEND

Live music & theater entertainment, comedy shows, a variety of  
activities, chocolate shops, and high-end restaurants make New  
Brunswick an excellent choice for couples on Valentine's Day!

It's easy to see why this city is romantic with so many fun date night  
options, luxurious spas, and delicious dining selections.

[2]Book Your Reservations Now

Jersey's Choice Restaurant Poll

Presented by [3]New Brunswick's Performing Arts Center, it's your time  
to

show your support for your favorite spots in New Brunswick as NJ

Monthly hosts their 39th Annual Restaurants of NJ Poll!

New Brunswick is represented strongly with many of your favorites  
including, [4]Roosterspin, [5]Stage Left Steak, [6]Harvest Moon  
Brewery, [7]The Frog & The Peach, [8]Delta's, [9]Cambo Box,

- +- 4/4344: @NBCityCenter Vote for New Brunswick's Favorite Restaurants



--00000000000013d53505d7114793

Content-Type: image/png; name="screenshot.png"

Content-Disposition: attachment;  
filename="screenshot.png"

Content-Transfer-Encoding: base64


X-Attachment-Id: f\_kz66ba9k0

Content-ID: <f\_kz66ba9k0>

iVBORw0KGgoAAAANSUhEUgAABSGAAAGCCAYAAADjWS0TAAAKC2LDQ1BJQ0MgUHJvZmlsZQAASImF  
lgdUFFkWhl9V50RqaHL00UMDkn0SHEWlaXKYaakICREVGFFEREARZIGKjkqQQUREQWUQUMCADsgg  
oIyDAVBReWuY0ZndPbu3zj3vq9vv/XXfqz6nfgBILgwWKxbmASauPontaw8l5R8QKIWbBihaBgAQ  
gS6DmciydHd3Qe7AX+M/Y2kUQKvjffVVrf/8/X8Gb2hYIhMAyB1h0pPFTkL4IMLeqUmsVR5HmJ+N  
NIXwwipHrDGMXuWQdRZam+PtaY2wGgB4MoPBjkCapyN1qRRmBKJD9EdYKz40Kh7hVX0zZiQjFOE+  
hNXCy5PTEH6/0icubhtSJ8khrBTyN82If+iHfNNnMCK+cVxsMvPPfa2eCDks3scLGUWQFAPhQAPE  
gmSQBqQAC7DBNqQShVTckLP/7+voa+uskZkssB1ZEQUIQCRIQtbb/U3La00pCaQCBjInDKm4IJf1  
6ntcl3xHW10FaHe+1zKRvZqGcjicju81F2TPL/cje3nxvabYCQCXIAB9ucxkdsp6bfXoAQb5b3AD  
fiAMJIAsUALqQAcYABNgAWyBE3AD3iAAbAFMpN84pKtUkAH2gmyQC46A46AE1I0zoBacBxdBK+gA  
18EtcBcMghHwBEyAafAKLIAlsAJBEA6iQFRIGJKE5CFVSAeiQ2aQLeQCeUIBUDAUAacVDyVAGtA/K  
h0aaEaaCaoN+aa5A16Hb0BD0C1aE5aC30CcYB7NhflacVoA1YTpsCTyD3vBm0A10aNPbLPawXaxX

# Mail Access Protocols

# Mail access protocols



- SMTP: delivery/storage to receiver's server. Focused on **push**
- Mail access protocol: retrieval from server
  - POP: Post Office Protocol [RFC 1939]
    - Client connects to POP3 server on TCP port 110
  - IMAP: Internet Mail Access Protocol [RFC 1730]
    - Client connects to TCP port 143
  - HTTP: gmail, outlook, etc.

# POP vs IMAP

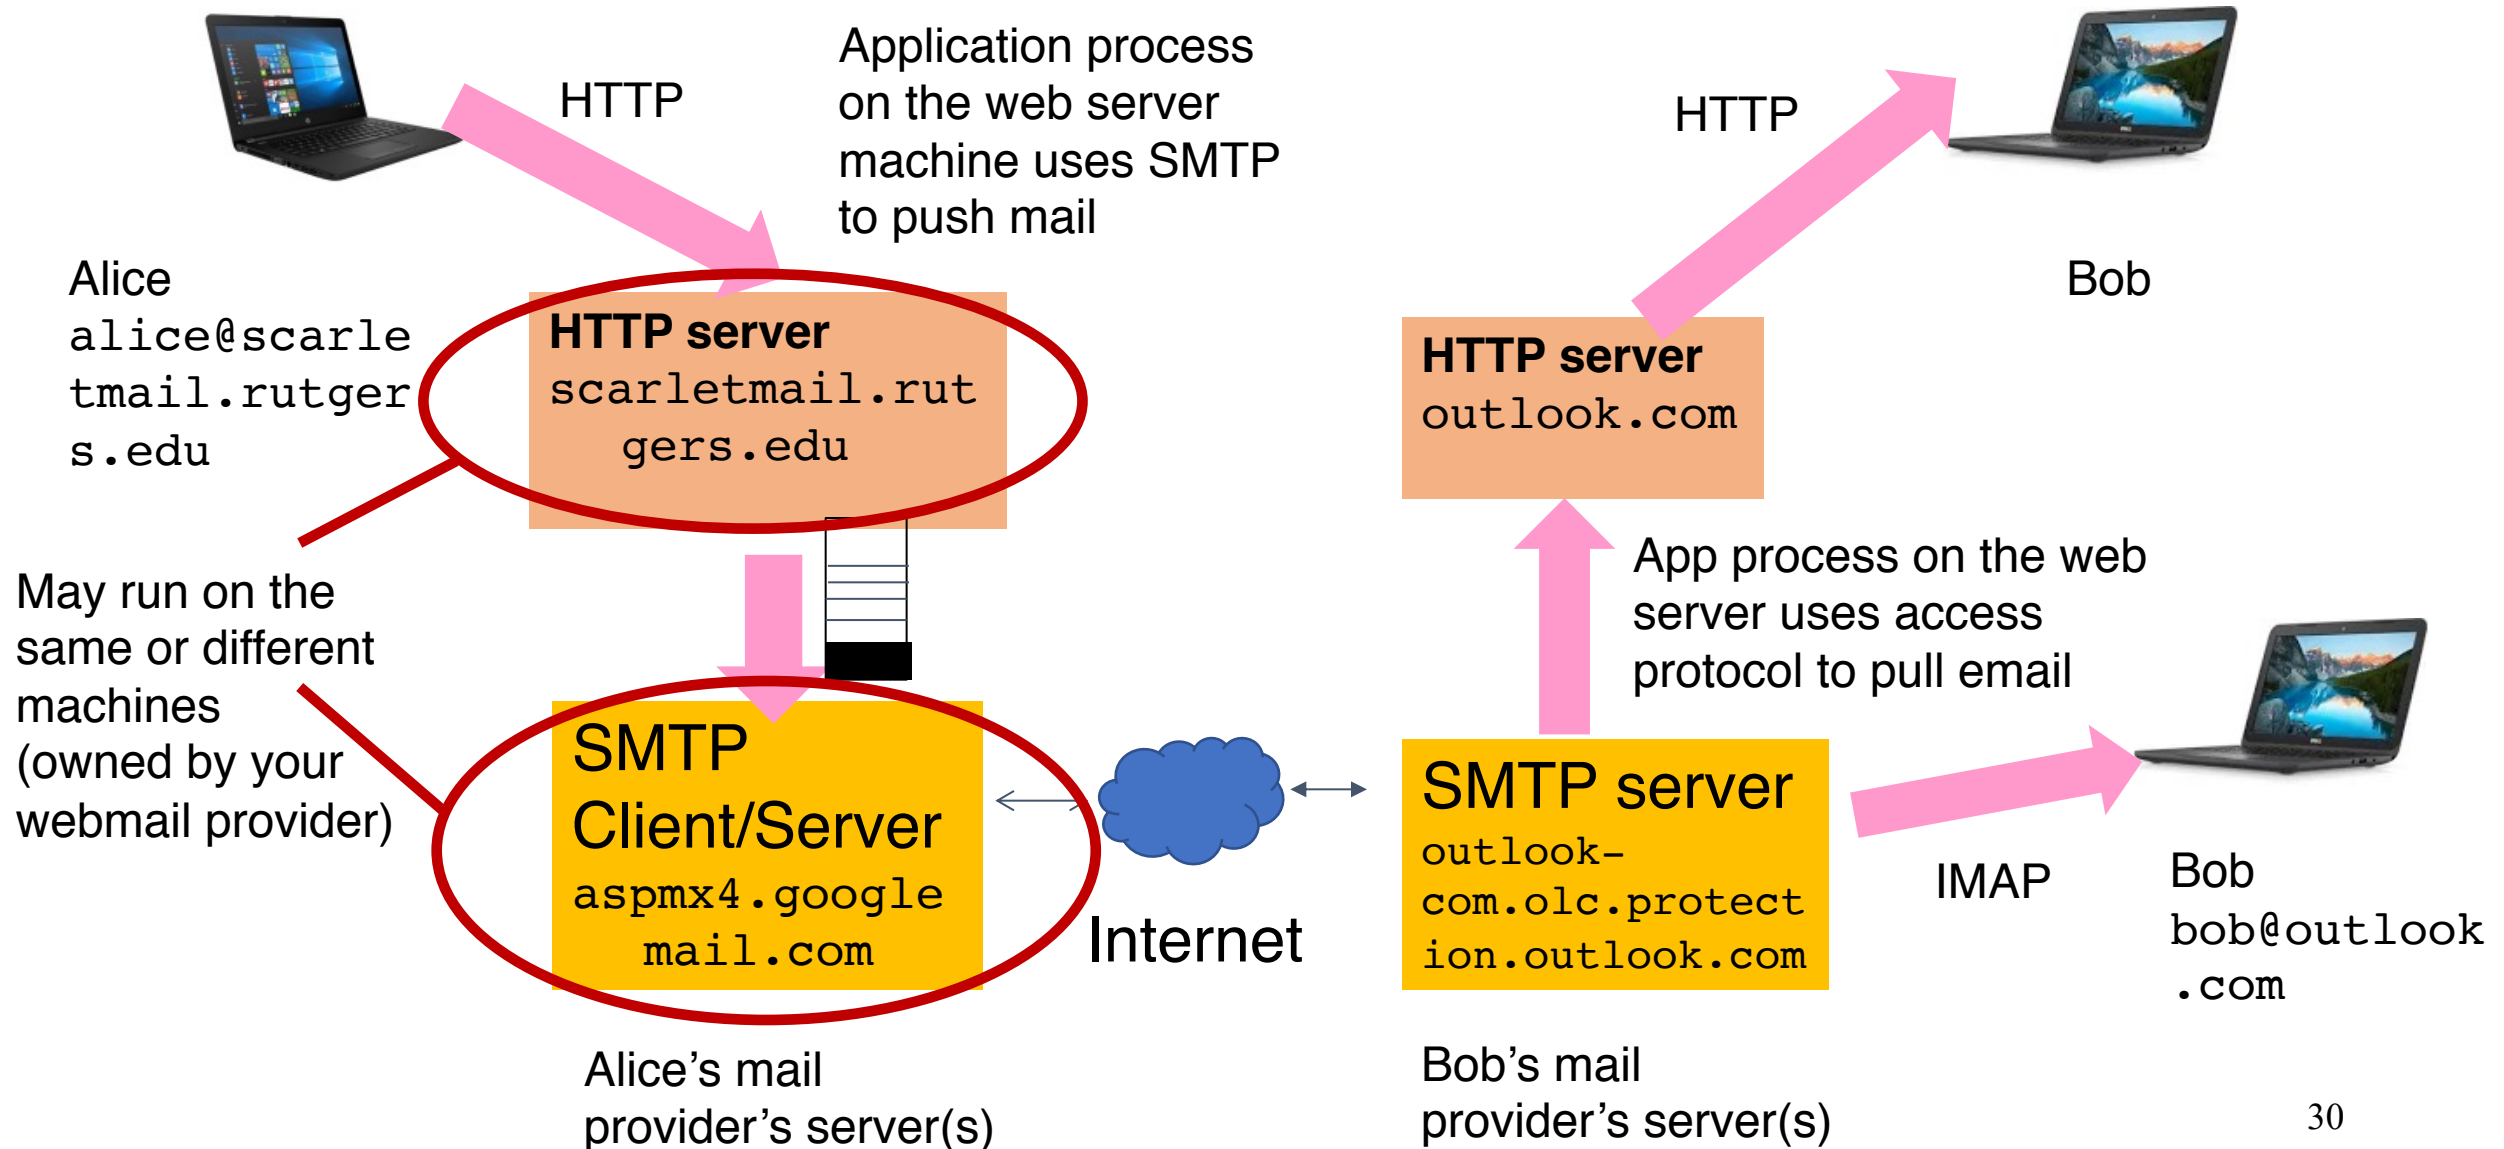
- POP3
- **Stateless** server
- UA-heavy processing
- UA retrieves email from server, then typically deleted from server
- Latest changes are at the UA
- Simple protocol (`list`, `retr`, `del` within a POP session)

- IMAP4
- **Stateful** server
- UA and server processing
- Server sees folders, etc. which are visible to UAs
- Latest changes are at the server
- Complex protocol
- Heavily used: email sync across devices, reliable, ...

# What about web-based email?

- Connect to mail servers via web browser
  - Ex: gmail, scarletmail, etc.
- Browsers speak HTTP
- Email servers speak SMTP
- Need to bridge these two

# Web based email



# Comparing SMTP with HTTP

- HTTP: pull
- SMTP: push
- both have ASCII command/response interaction, status codes
- HTTP: each object encapsulated in its own response msg
- SMTP: multiple objects sent in multipart msg
- HTTP: can put non-ASCII data directly in response (dedicated entity body for binary data)
- SMTP: need ASCII-based encoding (base64)

# More themes from app-layer protocols

- **Keep it simple until you really need complexity**
  - Start with ASCII-based design; stateless servers. Then introduce:
  - Cookies for HTTP state
  - Stateful mail (IMAP, folders, etc.) for email organization
  - Security extensions (e.g., TLS)
  - Performance optimizations: persistence, caching, indirection, ...
  - Use headers as much as possible to non-intrusively evolve functionality
- **Partition functions based on what's done best at the user (app) and protocol. Examples:**
  - Content rendering for users (browser, UA) separate from protocol operations (mail server)
  - mail UA doesn't need to be "always on" to send or receive email reliably. That's the mail server's job