CSC 102 • HOW THE INTERNET WORKS FINAL TAKE-HOME EXAM DUE 5:00 PM ON TUESDAY, 4 OCTOBER 2011 PLEASE SUBMIT YOUR WORK ELECTRONICALLY VIA MOODLE

You may not consult any person other than the professor when completing this exam, but you may use references online and in print. Wherever you use external sources, **please include a URL or citation with your answer**. Don't forget to use quotation marks when copying verbatim from a source. There will be a certain amount of Web searching necessary to respond to several of the questions. Wikipedia will be accepted as a source of answers for purposes of this exam. You may also have to use Google or its equivalent intelligently, and some of the other tools we have studied.

1. Data Transfer (12 points)

The Google N-Gram Corpus is a set of statistics of word sequence frequency extracted from analysis of over one trillion words of text documents. It has been made available for free to researchers and anybody else who would like to download it. The data come compressed in a 24 GB file.

a.) Suppose that you try to download it over a 56 kbps modem. How many hours will this take?

b.) Suppose that instead you have a friend mail you a copy of the data on 6 DVDs. If they take three days (72 hours) to arrive, what was the effective transfer rate?

c.) What connection speed (in Mb per second) would you need to have in order to download the entire corpus in three hours?

2. Communications Protocols (12 points)

Suppose that you have a simple home network with a wireless router connected to a cable modem acting as a bridge. Two laptops connect wirelessly to the router, and a desktop computer is plugged in using an Ethernet cable. Assume that the router is configured to act as a firewall, ignoring any unsolicited traffic from unknown addresses. Answer the following questions with specific reference to the protocols we have studied in class. (In other words, along with your answer, tell me the name of the protocol(s) you are basing it on!)

a.) The household connects to the Internet as IP address 102.221.87.165. Which devices in the home network see this address?

b.) When someone using a laptop visits a web site by clicking on a link, and the web server sends back a packet in response, how does the router know to pass it along to the laptop instead of blocking it at the firewall?

c.) Suppose that both the laptop and the desktop request a web page from www.smith.edu at the same time. Packets are sent to the router intended for both recipients, both addressed to 102.221.87.165. How does the router know which packet goes to which computer?

3. HTML (20 points)

Identify the best tag or escape code to use for each of the following tasks:

- a.) A list of web links
- b.) The column headings in a table
- c.) A separator line (XHTML style)
- d.) An ampersand symbol
- e.) A copyright symbol: ©
- f.) Associate a CSS file with a web page
- g.) Identify a major section of a page
- h.) Identify a subsection of a page
- i.) Identify text that should be typed on a keyboard

j.) Augment an image with "hot spots" that act as links. (We didn't talk about this one in class; you'll have to investigate it on your own. Think of this as testing your ability to extend your base knowledge.)

4. Images (12 points)

Explain why each of the following image types is not recommended for the application described, and recommend a better format to use.

a.) Windows bitmap (BMP) image for your company logo.

b.) TIFF image for a fancy link button image.

c.) JPEG image for a picture of a castle; the page background should show behind the castle towers.

5. HTML Forms (12 points)

Checkboxes, radio buttons, and popup menus each allow a user to make selections from multiple options. Describe the functional differences between the three controls, and explain the circumstances under which it is best to use each one.

6. Email (16 points)

Consider the email message source below. (Some irrelevant portions of the email have been redacted.) Please answer the questions that follow.

Return-path: <prvs=0077d39924=roberts.kendall@unionps.org> Received: from mscreen4.smith.edu ([131.229.64.70]) by gwsmtp1.smith.edu with ESMTP; Wed, 06 Apr 2011 11:31:13 -0400 Received: from ewsapp2.smith.edu (ewsapp1.smith.edu [131,229,64,90]) by mscreen4.smith.edu (8.14.3/8.14.3) with SMTP id p36FVD6c010253 for <nhowe@smith.edu>; Wed, 6 Apr 2011 10:31:13 -0500 Received: from (unknown [174.76.24.10]) by ewsapp2.smith.edu with smtp id 5f7d 000c e6f29578 6062 11e0 b131 001517ed064a; Wed, 06 Apr 2011 11:31:11 -0400 Authentication-Results: wg1.unionps.org header.from=roberts.kendall@unionps.org; domainkeys=neutral (no sig) Received: from mail.unionps.org (unknown [10.50.1.244]) by wg1.unionps.org (Extensible Content Security) with ESMTP id BDAB35017A9A90AC; Wed, 6 Apr 2011 10:31:07 -0500 (CDT) Received: from UnionEmail-Cert.unionps.org (10.50.1.243) by UnionAdminEmail.unionps.org (10.50.1.244) with Microsoft SMTP Server (TLS) id 8.2.176.0; Wed, 6 Apr 2011 10:31:11 -0500 Received: from UnionEmail-Cert.unionps.org ([fe80::a89e:c6c3:d509:f2ac]) by UnionEmail-Cert.unionps.org ([fe80::a89e:c6c3:d509:f2ac%10]) with mapi; Wed, 6 Apr 2011 10:31:10 -0500 From: "Roberts, Kendall" <roberts.kendall@unionps.org> To: "li@s.com" <li@s.com> Date: Wed, 6 Apr 2011 10:31:10 -0500 Subject: =?windows-1256?Q?Mailbox Alert!=FE?= Thread-Topic: =?windows-1256?Q?Mailbox_Alert!=FE?= Thread-Index: AQHL9G+n9NzS/CjvUkWq145NYHDeVg== Message-ID: <12719E9B2DDE40439DBB8DF1A3725DD9C3476026F7@UnionEmail-Cert.unionps.org> Content-Type: multipart/alternative; boundary=" 000 12719E9B2DDE40439DBB8DF1A3725DD9C3476026F7UnionEmailCer " MIME-Version: 1.0 Received-SPF: none --_000_12719E9B2DDE40439DBB8DF1A3725DD9C3476026F7UnionEmailCer_ Content-Type: text/plain; charset="windows-1256" Content-Transfer-Encoding: quoted-printable This is to notify you that you are over your mailbox limit which is 250MB=

This is to notify you that you are over your mailbox limit which is 250MB= as set by your mailbox manager, you are currently at 257MB, you will not= be able to create new e-mail to send or receive messages until you valid= ate your mailbox. To re-validate your account, click here:<http://gambiac= hat.com/form/use/IH/form1.html> Help Desk

--_000_12719E9B2DDE40439DBB8DF1A3725DD9C3476026F7UnionEmailCer_ Content-Type: text/html; charset="windows-1256" Content-Transfer-Encoding: quoted-printable

```
<html dir=3D"ltr"><head>
<meta http-equiv=3D"Content-Type" content=3D"text/html; charset=3Dwindows=</pre>
-1256">
<style title=3D"owaParaStyle"><!--P {
       MARGIN-TOP: 0px; MARGIN-BOTTOM: 0px
}
--></style>
</head>
<body ocsi=3D"x">
<div dir=3D"ltr"><font face=3D"Tahoma" color=3D"#000000" size=3D"2"><span=</pre>
class=3D"Apple-style-span" style=3D"WORD-SPACING: 0px; FONT: medium/20px=
'Times New Roman'; TEXT-TRANSFORM: none; COLOR: rgb(0,0,0); TEXT-INDENT:=
0px; WHITE-SPACE: normal; LETTER-SPACING: normal; BORDER-COLLAPSE: separ=
ate; orphans: 2; widows: 2; -webkit-border-horizontal-spacing: 0px; -webk=
it-border-vertical-spacing: 0px; -webkit-text-decorations-in-effect: none=
; -webkit-text-size-adjust: auto; -webkit-text-stroke-width: 0px"><span c=
lass=3D"Apple-style-span" style=3D"FONT-SIZE: small; LINE-HEIGHT: normal;=
FONT-FAMILY: Arial"><font style=3D"TEXT-INDENT: 0px! important; LINE-HEI=
GHT: normal" face=3D"Tahoma" color=3D"#000000" size=3D"2">This
is to notify you that you are over your mailbox limit which is 250MB as =
set by your mailbox manager, you are currently at 257MB, you will not be =
able to create new e-mail to send or receive messages until you validate =
your mailbox. To re-validate your account,<span class=3D"ecxApple-convert=
ed-space" style=3D"TEXT-INDENT: 0px! important; LINE-HEIGHT: 17px"> =
c/span><font style=3D"TEXT-INDENT: 0px! important: LINE-HEIGHT: normal" c=
olor=3D"#0066cc"><font class=3D"ecxApple-style-span" style=3D"TEXT-INDENT=
: 0px! important; LINE-HEIGHT: normal" color=3D"#0068cf"><span class=3D"e=
cxApple-style-span" style=3D"CURSOR: default: TEXT-INDENT: 0px! important=
; LINE-HEIGHT: 17px"><u style=3D"TEXT-INDENT: 0px! important; LINE-HEIGHT=
: 17px"><a style=3D"CURSOR: default; COLOR: rgb(0,104,207); TEXT-INDENT: =
0px! important; LINE-HEIGHT: 17px; TEXT-DECORATION: underline" onclick=3D=
"onClickUnsafeLink(event);" href=3D"http://gambiachat.com/form/use/IH/for=
m1.html" target=3D"_blank">click
here:</a></u></span></font></font>&nbsp;</font>
<div style=3D"TEXT-INDENT: 0px! important; LINE-HEIGHT: 17px"><font style=</pre>
=3D"TEXT-INDENT: 0px! important: LINE-HEIGHT: normal" face=3D"Tahoma" col=
or=3D"#000000" size=3D"2">Help Desk</font></div>
</span></span></font></div>
</bodv>
</html>
```

--_000_12719E9B2DDE40439DBB8DF1A3725DD9C3476026F7UnionEmailCer_--

a.) What is the final line of the email header in the message source as shown?

b.) This message uses the MIME format. How many parts are there?

c.) The text of this message is obviously suspicious. Aside from that, what factor(s) about the email header are also suspicious?

d.) This email links to a web site. If the HTML version of this message is viewed, what text will appear on the link? To what URL does the link actually lead?

7. Certificates (8 points)

Suppose that you visit a web site of a merchant (Zed's Emporium) who sells a product you wish to buy. The browser displays a green bar next to the address, and closer inspection reveals that the web site has a Class 3 certificate registered to Zed Raskell Mercantile, Inc. by VeriSign, and that communication with the site is protected by 128-bit encryption. What guarantees of security does this provide? What are the risks (if any) of providing a credit card number to an order form on the web site?

8. Web Search (8 points)

Consider the three scenarios below. Assume all have the same links from external sites pointing to the home page, and no external links into any of the subpages. Rank the three configurations according to the page rank of the home page – in which case will it be the highest? Lowest? Explain why.

