UNDERSTANDING PHISHING

INKY

Hidden Text and Zero Font Attacks

You may have heard that some phishing emails use something called zero font — text hidden in an email by various means. But what exactly is this technique, and why do attackers use it? In this document we'll explain both the why and how of hidden text and zero font.

Hidden Text — What the Font?

You may have heard that some phishing emails use something called *zero font* — text hidden in an email by various means. But what exactly is this technique, and why do attackers use it? In this document we'll explain both the why and how of hidden text and zero font.

Zero font is a narrow term for a specific way bad actors can hide text in an email, but you can think of this exploit more broadly as *hidden text*. Most email is now designed using HTML, the same markup language used to render web pages. And while email itself has been around since 1971, the so-called MIME standards that permit HTML and other kinds of rich markup in email only date to the mid 90s — relatively new for email.

By embracing HTML email, the standards bodies paved the way for highly styled electronic messages that recipients could directly engage with in myriad ways. But along with the intrinsic utility of styled email came various security implications as well; hidden text is just one of many.

The richness and complexity of modern HTML make it very difficult to properly interpret. Browser engines that render HTML contain many millions of lines of code contributed by thousands of developers over several decades. While we may think of web pages as being simple, they are anything but. And with HTML emails now pervasive, that complexity has crept into our email systems.

Why does this complexity matter?

Because now it's very hard for software to determine what the end user will actually see when confronted with a given email. And this gives attackers new ways to slip malicious content through mail protection systems.

Consider this zero-day phishing email that INKY blocked, displayed in Chrome:

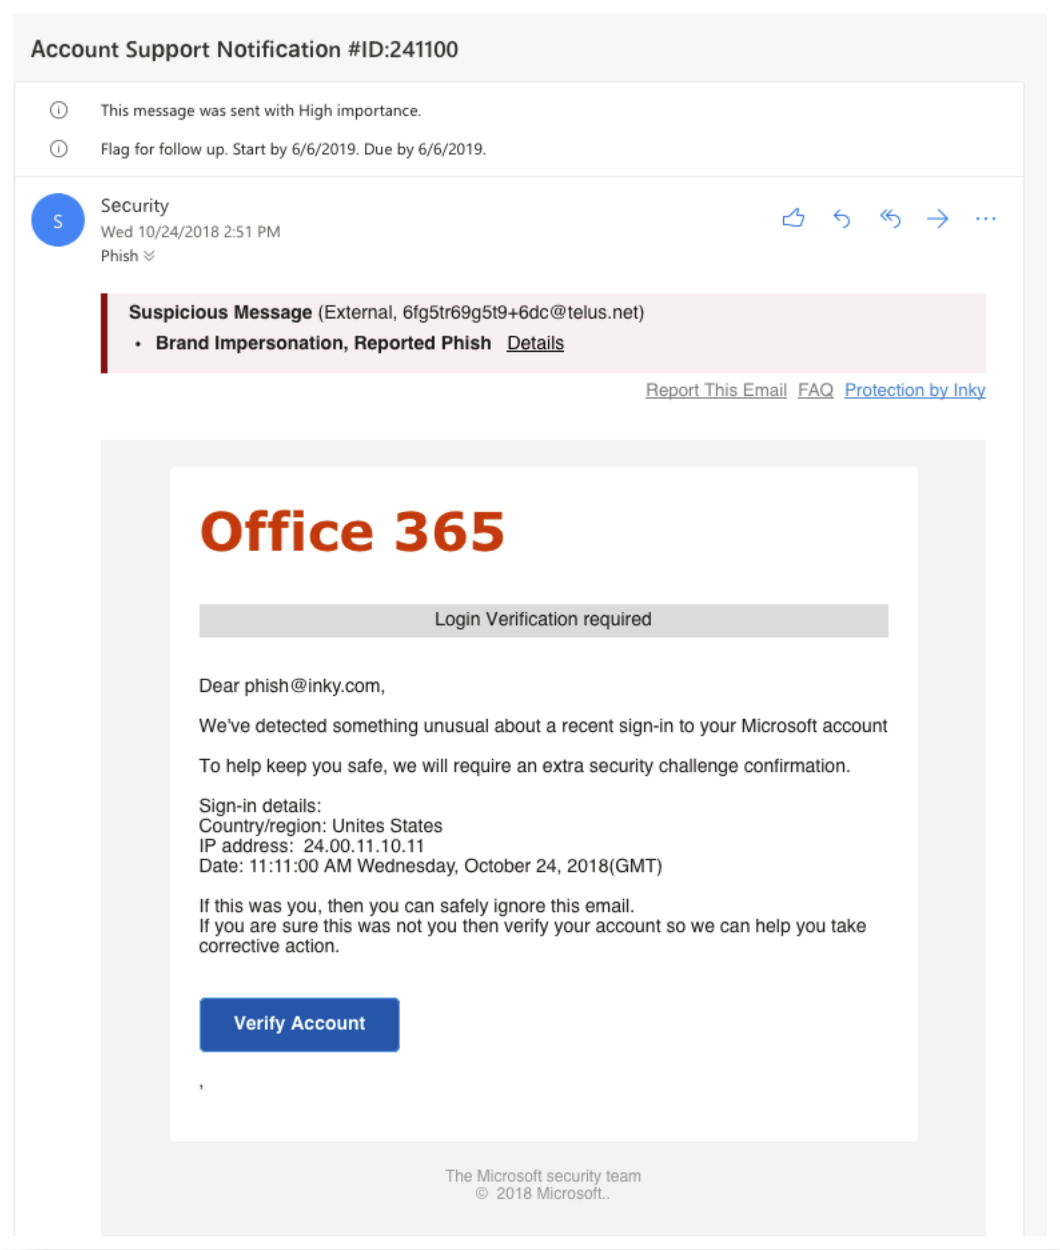

Account Support Notification #ID:241100

This message was sent with High importance.

Flag for follow up. Start by 6/6/2019. Due by 6/6/2019.

Security
Wed 10/24/2018 2:51 PM
Phish

Suspicious Message (External, 6fg5tr69g5t9+6dc@telus.net)

- **Brand Impersonation, Reported Phish** Details

Report This Email FAQ Protection by Inky

Office 365

Login Verification required

Dear phish@inky.com,

We've detected something unusual about a recent sign-in to your Microsoft account

To help keep you safe, we will require an extra security challenge confirmation.

Sign-in details:
Country/region: Unites States
IP address: 24.00.11.10.11
Date: 11:11:00 AM Wednesday, October 24, 2018(GMT)

If this was you, then you can safely ignore this email.
If you are sure this was not you then verify your account so we can help you take corrective action.

,

The Microsoft security team
© 2018 Microsoft..

This mail is a brand forgery email intended to harvest login credentials. Upon clicking *Verify Account* the user is taken to what looks like a normal Office 365 login page. If she enters her email address and password into this page, however, she's just given this information to the scammers, who will then send phishing emails from her account (and as a bonus, try this same password on other sites as well). Now take a look at this same mail in Microsoft Outlook for Windows:

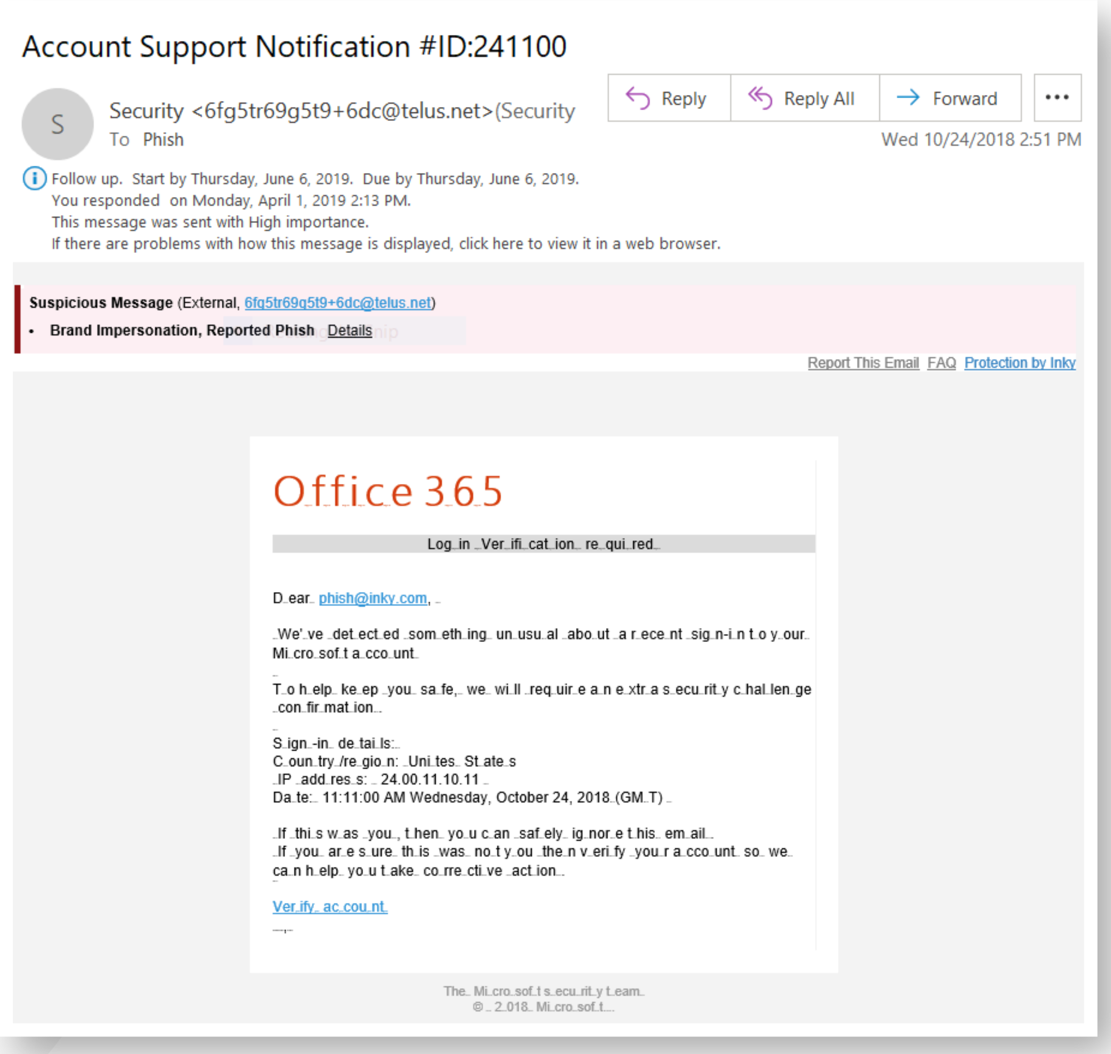

Account Support Notification #ID:241100

Security <6fg5tr69g5t9+6dc@telus.net>(Security
To Phish
Wed 10/24/2018 2:51 PM

Follow up. Start by Thursday, June 6, 2019. Due by Thursday, June 6, 2019.
You responded on Monday, April 1, 2019 2:13 PM.
This message was sent with High importance.
If there are problems with how this message is displayed, click here to view it in a web browser.

Suspicious Message (External, 6fg5tr69g5t9+6dc@telus.net)

- **Brand Impersonation, Reported Phish** Details

Report This Email FAQ Protection by Inky

Office 365

Log_in _Ver_ifi_cat_ion_ re_qui_red_

D_ear_ phish@inky.com, _

_We'_ve _det_ect_ed _som_eth_ing_ un_usu_al _abo_ut _a r_ece_nt _sig_n-i_n t_o y_our_ Mi_cro_sof_t a_cco_unt_

_

T_o h_elp_ ke_ep _you_ sa_fe,_ we_ wi_ll _req_uir_e a_n e_xtr_a s_ecu_rit_y c_hal_len_ge _con_fir_mat_ion_

_

S_ign_-in_ de_tai_ls:_
C_oun_try_/re_gio_n: _Uni_tes_ St_ate_s
_IP _add_res_s: _ 24.00.11.10.11 _
Da_te:_ 11:11:00 AM Wednesday, October 24, 2018_(GM_T) _

_If _thi_s w_as _you_, t_hen_ yo_u c_an _saf_ely_ ig_nor_e t_his_ em_ail_
_If _you_ ar_e s_ure_ th_is _was_ no_t y_ou _the_n v_eri_fy _you_r a_cco_unt_ so_ we_ ca_n h_elp_ yo_u t_ake_ co_rre_cti_ve _act_ion_

Ver_ify_ ac_cou_nt_

The_ Mi_cro_sof_t s_ecu_rit_y t_eam_
© _ 2_018_ Mi_cro_sof_t_

See how different it looks? That’s because Outlook has its own HTML rendering engine — one that differs from the ones used in mainstream browsers.

This Windows Outlook quirk gives us visibility into what’s going on here. Notice that in Outlook the big red Office 365 text looks like it has gaps between the letters? That’s zero font at work! Let’s see what’s going on in the HTML:

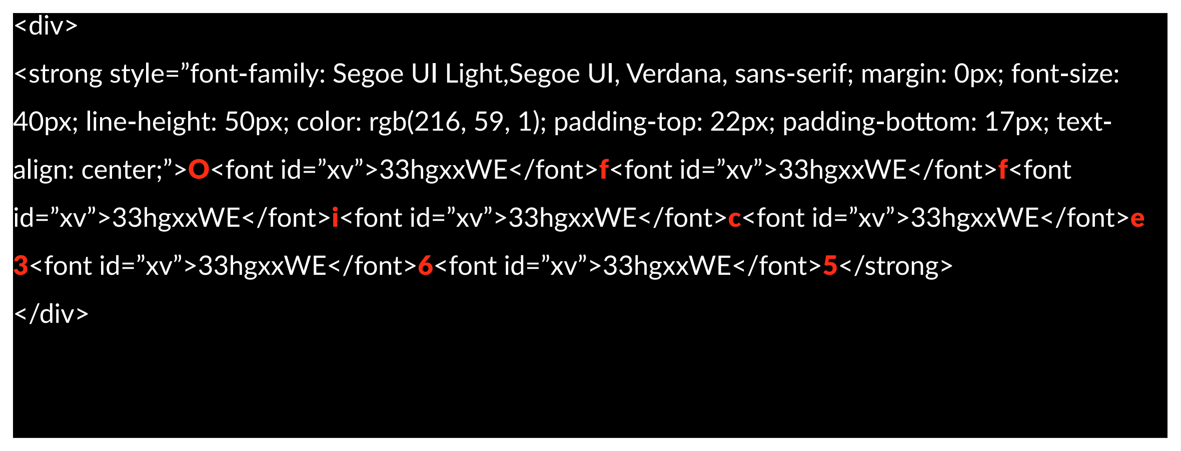

```
<div>
<strong style=”font-family: Segoe UI Light,Segoe UI, Verdana, sans-serif; margin: 0px; font-size:
40px; line-height: 50px; color: rgb(216, 59, 1); padding-top: 22px; padding-bottom: 17px; text-
align: center;”>O<font id=”xv”>33hgxxWE</font>f<font id=”xv”>33hgxxWE</font>f<font
id=”xv”>33hgxxWE</font>i<font id=”xv”>33hgxxWE</font>c<font id=”xv”>33hgxxWE</font>e
3<font id=”xv”>33hgxxWE</font>6<font id=”xv”>33hgxxWE</font>5</strong>
</div>
```

The red highlights spell out Office 365; this is the big red logotype in the upper left. But what’s all that other text in there? Let’s highlight in yellow:

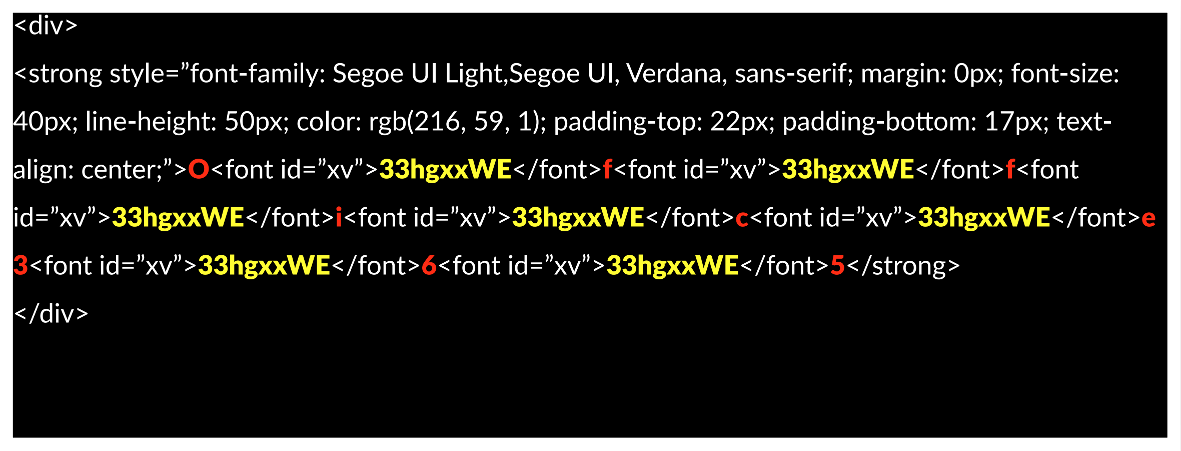

```
<div>
<strong style=”font-family: Segoe UI Light,Segoe UI, Verdana, sans-serif; margin: 0px; font-size:
40px; line-height: 50px; color: rgb(216, 59, 1); padding-top: 22px; padding-bottom: 17px; text-
align: center;”>O<font id=”xv”>33hgxxWE</font>f<font id=”xv”>33hgxxWE</font>f<font
id=”xv”>33hgxxWE</font>i<font id=”xv”>33hgxxWE</font>c<font id=”xv”>33hgxxWE</font>e
3<font id=”xv”>33hgxxWE</font>6<font id=”xv”>33hgxxWE</font>5</strong>
</div>
```

We can see that the attacker has inserted the gibberish text “33hgxxWE” in between the logotype letters. And when we look at the CSS styles, we can see this:

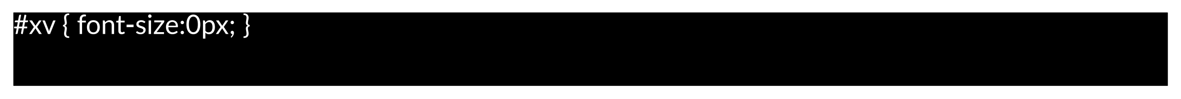

```
#xv { font-size:0px; }
```

The size of font “xz” (the yellow text) is set to zero, making the gibberish invisible! This is in fact where the name zero font comes from. So why do they this? To hide from mail protection software. This text is invisible to the end user, but it makes the word “Office” in the logotype look like

O33hgxxWEf33hgxxWEf33hgxxWEi33hgxxWEc33hgxxWEe

to the mail protection software. So if the software is looking for brand-indicative text like “Office 365”, it won’t find a match. This tactic therefore prevents legacy mail protection systems from classifying this mail as appearing to be from Microsoft. And since it doesn’t know it appears to be from Microsoft, it doesn’t require the mail to be from a Microsoft-controlled mail server. So it sails right through, ending up in the victim’s inbox. And if the victim is using a standard browser to read the email, she’ll just see “Office 365” — and potentially fall for the scam.

Attackers can embed text into their emails that is both invisible to end users and visible — and confusing — to the machines that automatically scan the mail looking for signs of malicious intent or branding.

Take this Phish and Stuff it!

Another trick attackers use is *keyword stuffing*. This tactic relies on hidden text too. Let's look at an example from a fake Bank of America email:

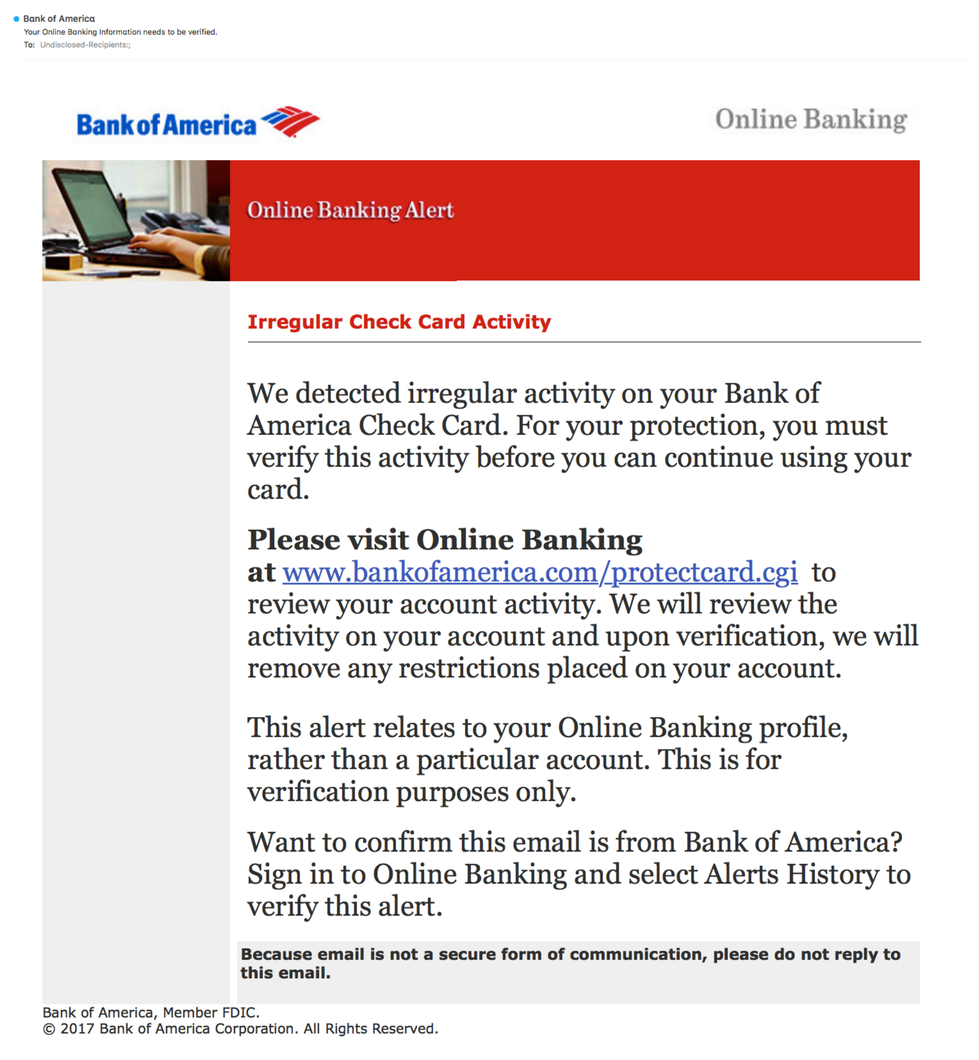

Bank of America
Your Online Banking Information needs to be verified.
To: Undisclosed-Recipients:;

Bank of America

Online Banking

Online Banking Alert

Irregular Check Card Activity

We detected irregular activity on your Bank of America Check Card. For your protection, you must verify this activity before you can continue using your card.

Please visit Online Banking at www.bankofamerica.com/protectcard.cgi to review your account activity. We will review the activity on your account and upon verification, we will remove any restrictions placed on your account.

This alert relates to your Online Banking profile, rather than a particular account. This is for verification purposes only.

Want to confirm this email is from Bank of America? Sign in to Online Banking and select Alerts History to verify this alert.

Because email is not a secure form of communication, please do not reply to this email.

Bank of America, Member FDIC.
© 2017 Bank of America Corporation. All Rights Reserved.

Looks pretty normal, right? But this email, too, contains text invisible to the end user meant to confuse the mail protection system.

Here’s the same mail with a tiny modification to the HTML to make white-on-white text in this email display as red-on-white:

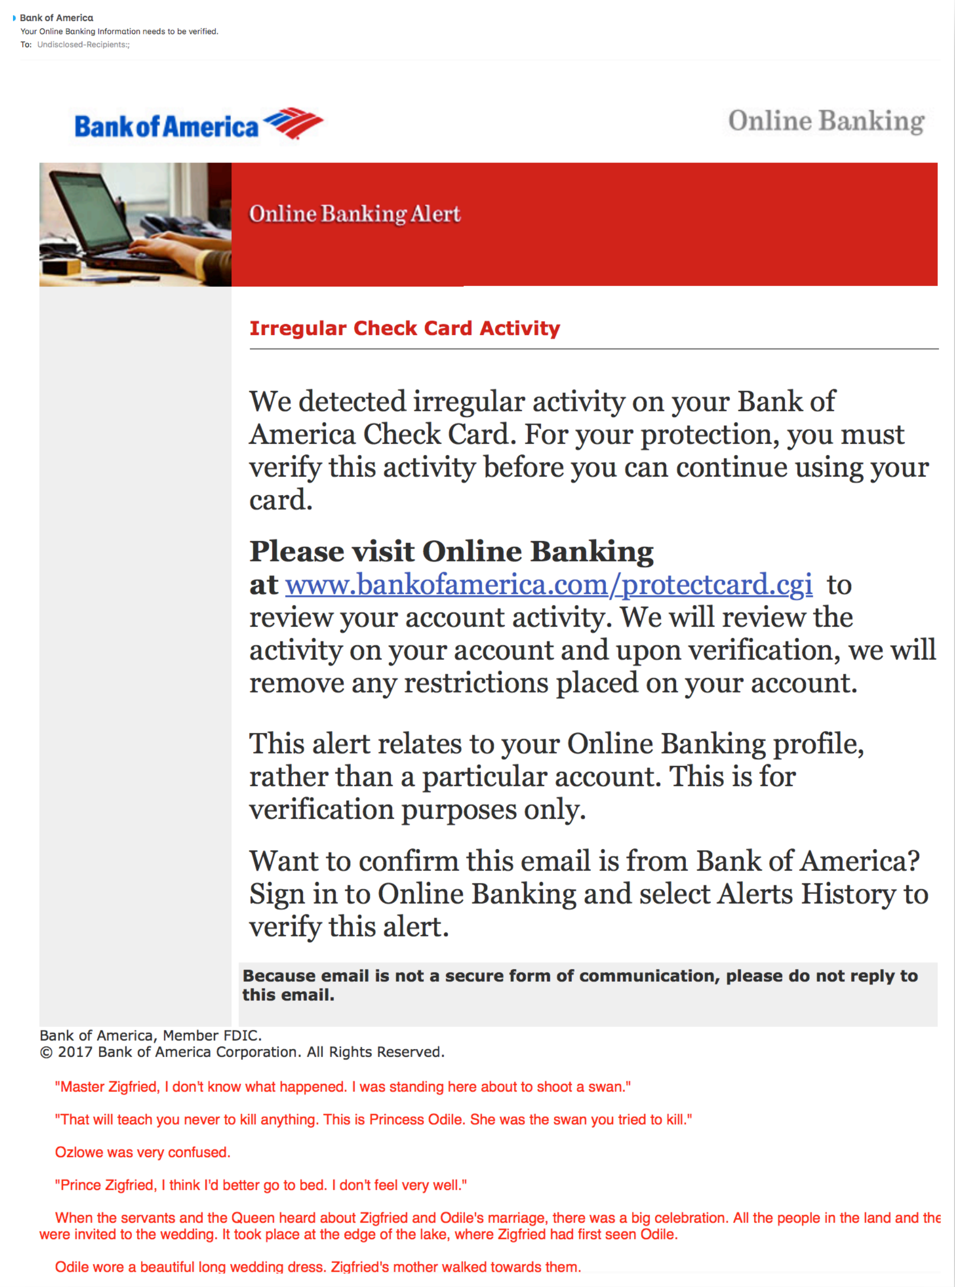

Bank of America
Your Online Banking Information needs to be verified.
To: Undisclosed-Recipients:;

Bank of America

Online Banking

Online Banking Alert

Irregular Check Card Activity

We detected irregular activity on your Bank of America Check Card. For your protection, you must verify this activity before you can continue using your card.

Please visit Online Banking at www.bankofamerica.com/protectcard.cgi to review your account activity. We will review the activity on your account and upon verification, we will remove any restrictions placed on your account.

This alert relates to your Online Banking profile, rather than a particular account. This is for verification purposes only.

Want to confirm this email is from Bank of America? Sign in to Online Banking and select Alerts History to verify this alert.

Because email is not a secure form of communication, please do not reply to this email.

Bank of America, Member FDIC.
© 2017 Bank of America Corporation. All Rights Reserved.

"Master Zigfried, I don't know what happened. I was standing here about to shoot a swan."

"That will teach you never to kill anything. This is Princess Odile. She was the swan you tried to kill."

Ozlowe was very confused.

"Prince Zigfried, I think I'd better go to bed. I don't feel very well."

When the servants and the Queen heard about Zigfried and Odile's marriage, there was a big celebration. All the people in the land and the were invited to the wedding. It took place at the edge of the lake, where Zigfried had first seen Odile.

Odile wore a beautiful long wedding dress. Zigfried's mother walked towards them.

All that red text is keyword stuffing! The attackers have put text at the end of the mail to confuse the mail protection system into thinking this is a conversational email rather than a transactional one apparently from well-known brand.

Attackers can also make text invisible to end users by setting its color to white-on-white. They will sometimes add text scraped from the web to an email to make the email appear to be a conversation between two people rather than a transactional email that might be a brand forgery.

The bottom line is that HTML email gives scammers a nearly endless supply of ways to hide text from end users, but still ensure the mail protection system sees it — and becomes confused by it.

How INKY Solves it

The richness of HTML — that there so many ways attackers can hide text — makes this a hard problem. What we need to do is to make the software see the email as closely as possible to how the end user will see it — it almost seems like we need to render the email just as though we were going to display it on a screen.

And that's exactly what INKY does.

For every HTML part of every email, INKY uses a standards-compliant rendering engine to create a pixel-perfect representation of the body of the email, just as the recipient will see it. INKY then applies a variety of computer vision and approximate matching techniques to "look" for both brand-indicative and scam-indicative visuals. All in under a second per email.

INKY can thus ensure that if the recipient would perceive the mail as being from a particular brand, the mail really is from a mail server controlled by that brand. This very general solution protects users from *any kind of hidden text attack* — even one INKY has never encountered before.

Conclusion

Phishing scams continue to evolve, and attackers have an ever growing bag of tricks with which to fool both human recipients and mail protection software. One of these tricks is to hide text within the HTML body of an email, such that while the end user doesn't see it, it confuses legacy mail protection systems.

INKY renders HTML email just like a real browser would, and then uses computer vision techniques to "see" the result much like a person would. This gives INKY a powerful tool to thwart entire class of hidden text and zero font attacks — even variants hackers haven't yet invented.