printout

Keystone MacCentral Macintosh Users Group ❖ www.keystonemac.com

Keystone MacCentral
February Meeting

Please see your membership email for the links
to this month's Zoom meeting or email us
at KeystoneMacCentral@mac.com.

During our program this month we plan to discuss

• OpenCore Legacy Patcher

Ventura drops support for the last of the pre-Retina MacBook Airs. OpenCore Legacy Patcher can add it back.

• ChatGTP

The tool lets you type natural-language prompts. ChatGPT offers conversational, if somewhat stilted, responses. It derives its answers from huge volumes of information on the internet.

For example, you can ask it encyclopedia questions like, "Explain Newton's laws of motion." You can tell it, "Write me a poem," and when it does, say, "Now make it more exciting." You can ask it to write a computer program that'll show you all the different ways you can arrange the letters of a word. It also works with other art forms — submit a portrait and ask it to render it in the style of your favorite artists. 🖤

We have virtual meetings via Zoom
on the third Tuesday of each month.

Emails will be sent out prior to each meeting.
Follow the directions/invitation each month
on our email — that is, just click on the link
to join our meeting.

Contents

By Adam Engst

Use Live Text to Digitize Your Cookbooks

I love cookbooks. I'm a sucker for paging through them and trying to imagine how difficult recipes will be and what they will taste like. Some authors, like J. Kenji López-Alt and Deb Perelman, are tremendously amusing, and cookbooks often have luscious photos that are always prettier than what I end up plating.

But my no-longer-secret shame is that after an early infatuation with a cookbook, it often ends up on my shelf, brought down only occasionally for a handful of recipes. Sometimes those favorites are marked with sticky tabs or bookmarks; more frequently, I resort to the index. Much as I approve of the Eat Your Books searchable recipe index site, I never managed to work it into my habits (see "Use the Web to Cook Your Books," 17 March 2022), so I often have to flip through a cookbook to find the one recipe I make repeatedly.

A few years ago, when Tristan was starting to cook on his own in college and asking for recipes for the foods he had grown up eating, we went all in on Paprika, a brilliant recipe app available for the iPhone, iPad, and Mac, along with Android and Windows (for our original coverage, see "FunBITS: Paprika Recipe Manager for iPhone, iPad, and Mac," 14 March 2014). I can't remember if there was a bundle deal then or not—the apps are now sold separately—but it's a perfect use of Family Sharing since Tonya, Tristan, and I can all now access our family recipes from whatever device we have handy.

The hurdle with Paprika—and any digital alternative to analog cookbooks—is importing recipes. It does a solid job of importing recipes from websites using the systemwide sharing extension or while viewing a site within Paprika's built-in Web browser. But most cookbooks don't have compani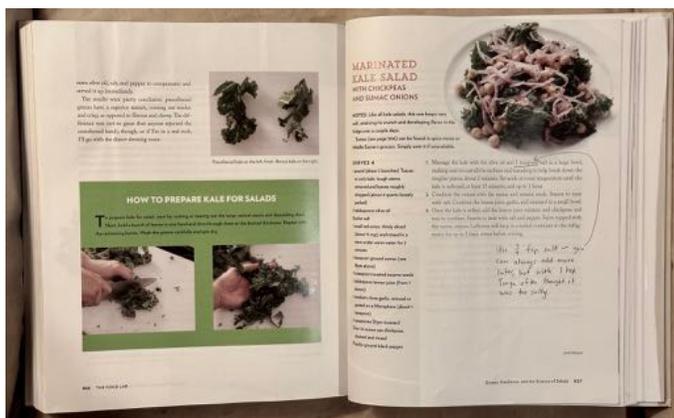on websites, and even cooking magazines like Cook's Illustrated often charge an annoying extra fee for digital access to the recipes you can read on paper.

My initial method of speeding up the process of transferring a recipe from a cookbook or magazine to Paprika was to use Voice Control's dictation, which is still more capable than the improved dictation available in the iOS/iPadOS 16 keyboard (see "How iOS and macOS Dictation Can Learn from Voice Control's Dictation," 31 August 2020). It's easy to read a recipe clearly enough to get a good transcription, although you must be careful to speak punctuation marks and line breaks. Between Web imports and dictated recipes, we now have over 200 favorites in Paprika, but I still frequently need to find recipes in our cookbooks.

With the rise of AI-driven image scanning, I recently wondered if Paprika would let me take a picture of a recipe, do OCR on the text, and recognize the layout to split out the description, ingredients, directions, and notes. Alas, the answer turns out to be no, and although such a feature does exist in other recipe apps, including CookBook, the reader-recommended Mela, and Recipe Keeper, I don't want to switch away from Paprika just for that feature.

However, once I was thinking along the lines of scanning, I realized that the Live Text feature Apple introduced in iOS 15 can insert text detected in the camera viewfinder. Although Live Text is unquestionably cool, I hadn't found any meaningful use for it until now. With Paprika, though, importing a recipe from a cookbook takes just a few minutes, far less time than dictating it.

Here's how I used Live Text to scan the "Marinated Kale Salad with Chickpeas and Sumac Onions" recipe from J. Kenji López-Alt's *The Food Lab*, complete with some handwritten notes.

Start by creating a new recipe in Paprika. Then follow these steps, which you can see illustrated in the screenshots below:

1. Tap the Name text field as though you were going to type in it.

2. Tap the Scan Text button that appears. The Live Text camera viewfinder replaces the keyboard at the bottom of the screen.

3. Position the text you want to scan in the viewfinder and verify that it has yellow marks around it.

4. Tap the blue insert button. (And no, I cannot fathom why Apple made it lowercase.)

5. Repeat with the rest of the fields in the recipe: Description, Ingredients, Directions, and more.

5.

A few comments and tips prompted by the screenshots:

In this recipe, the name is broken into multiple short lines, which works well with Live Text. When you're faced with a long line length, whether for the name or other parts of the recipe, it's difficult to get the desired text to fit in the viewfinder without also capturing more text above or below. You may have to insert more text than you want and delete it afterward.
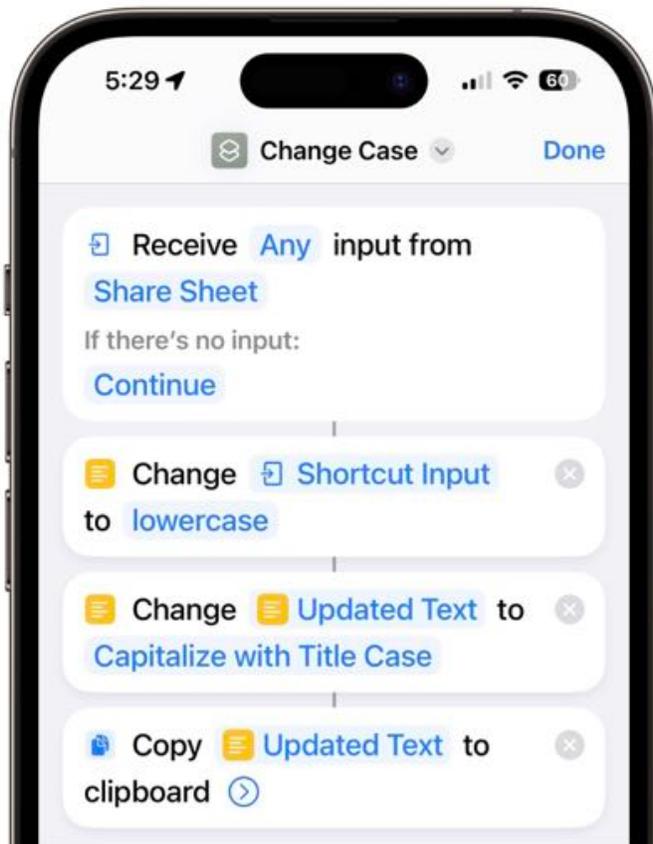
• Recipe names are often in all caps, which Live Text preserves. That bothered me, so I created a shortcut that converts any selected text to title case. Alas, it's a bit fussy to use. To change an all-caps name after inserting it, you must tap in the Name field, tap Select All, tap the share icon, select this Change Case shortcut, close the share sheet, tap the still-selected text, and tap Paste. The $2.99 Text Case app might eliminate the need to close the share sheet, but that's not a huge win. A tip of the hat to TidBITS Talk reader Nalarider, who pointed out that those with the Mac version of Paprika could use Keyboard Maestro or another Mac utility.

• Live Text previews the text to be inserted in the destination field, which can make you think that it has completed the task. But no, it's not done until you tap the blue Insert button to fix the previewed text in place. I still occasionally make that mistake.

• Line breaks are tough for Live Text. Sometimes it adds unnecessary line breaks; at other times, it runs text together that should be broken across lines. It's not terrible, but you'll often need to do some editing to get the visual spacing you want.

• Ingredient lists often contain proper fractions like ½, and I was impressed that Live Text captured and inserted them correctly. Paprika includes special keyboard shortcuts for inserting common fractions, but dictation only inserts regular characters, like 1/2.

• It's not uncommon for an ingredient list to be broken across two columns. When that happens, scan the left column first and then the right. Scanning both at the same time won't produce the

results you want. After you scan the left column, tap below it in the Ingredients field, and you'll notice that the Scan Text button has shrunk to just an icon. It works the same, though.

- Although Live Text has more trouble with handwriting than printed text, it still does well enough that it's usually worth accepting the scanned text and editing it later.

- Read through the entire recipe after scanning to make sure Live Text didn't make any errors that will cause serious consternation, like interpreting a scrawled "¼ teaspoon salt" as "4 teaspoon salt." Confusing "t" and "T" would also be bad. I haven't noticed any errors along these lines in the recipes I've imported, but some are inevitable, particularly with handwritten notes.

Obviously, all that's special about scanning recipe text is that many people have large collections of recipes on paper. Live Text will work its magic on any analog text you wish to digitize, whether into a dedicated app like Paprika or any app that accepts text input.

- Live Text previews the text to be inserted in the destination field, which can make you think that it has completed the task. But no, it's not done until you tap the blue Insert button to fix the previewed text in place. I still occasionally make that mistake.

- Line breaks are tough for Live Text. Sometimes it adds unnecessary line breaks; at other times, it runs text together that should be broken across lines. It's not terrible, but you'll often need to do some editing to get the visual spacing you want.

- Ingredient lists often contain proper fractions like ½, and I was impressed that Live Text captured and inserted them correctly. Paprika includes special keyboard shortcuts for inserting common fractions, but dictation only inserts regular characters, like 1/2.

- It's not uncommon for an ingredient list to be broken across two columns. When that happens, scan the left column first and then the right. Scanning both at the same time won't produce the results you want. After you scan the left column, tap below it in the Ingredients field, and you'll notice that the Scan Text button has shrunk to just an icon. It works the same, though.

- Although Live Text has more trouble with handwriting than printed text, it still does well enough that it's usually worth accepting the scanned text and editing it later.

- Read through the entire recipe after scanning to make sure Live Text didn't make any errors that will cause serious consternation, like interpreting a scrawled "¼ teaspoon salt" as "4 teaspoon salt." Confusing "t" and "T" would also be bad. I haven't noticed any errors along these lines in the recipes I've imported, but some are inevitable, particularly with handwritten notes.

Obviously, all that's special about scanning recipe text is that many people have large collections of recipes on paper. Live Text will work its magic on any analog text you wish to digitize. 🗑️

Hardware security keys

Hardware security keys offer the most secure protection you can use to lock down your online accounts.

Yubico's basic YubiKey security keys cost $25 for a USB-A model and $29 for a USB-C model. Both also support NFC wireless connections.

Apple now lets you protect your Apple ID and iCloud account with hardware security keys, a significant upgrade for those who want maximum protection from hackers, identity thieves, or snoops.

Hardware security keys are small physical devices that communicate with USB or Lightning ports or with NFC wireless data connections when you're logging on to a device or in to an account.

As of iOS 16.3 and MacOS 13.2, Apple published details on how to use security keys with iPhones, iPads and Macs.

By Adam Engst
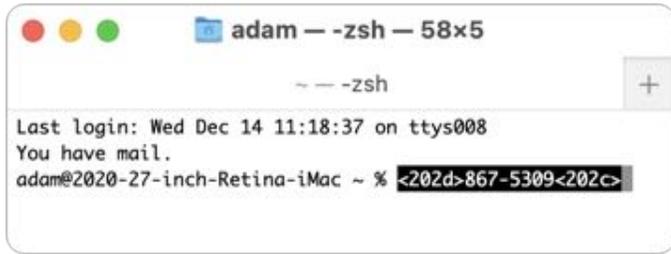
Beware When Pasting Phone Numbers Copied from Contacts

Have you ever had trouble pasting a phone number copied from Contacts into a Web form? It may not be a common problem, but it's one that TidBITS reader Dave Fultz solved after a frustrating tussle with an online order form.

Dave was trying to order a package in France, and the form requested a mobile number at the delivery address. To ensure he had it right, he copied it from Contacts and pasted it into the form, but when he clicked Submit, the form failed verification. He tried different formats for the phone number—in France, the leading zero in the ten-digit number should be included when calling inside the country but omitted when calling from abroad—to no effect. Pasting the number into a few online

verification services failed as well. Nor did pasting as plain text make any difference.

So Dave texted his friend at the number, which worked fine, and asked her to send him the number exactly as she would give it to a French person. It looked the same as his version, but when he copied and pasted the number from the Messages thread, the form accepted it! What was going on?

Dave was planning to use a Python script in Terminal to compare the number he had from Contacts with the one his friend had sent, but when he pasted his number into Terminal, the problem revealed itself instantly, as you can see in the highlighted text in my test below, which shows extra characters on either side of the phone number.
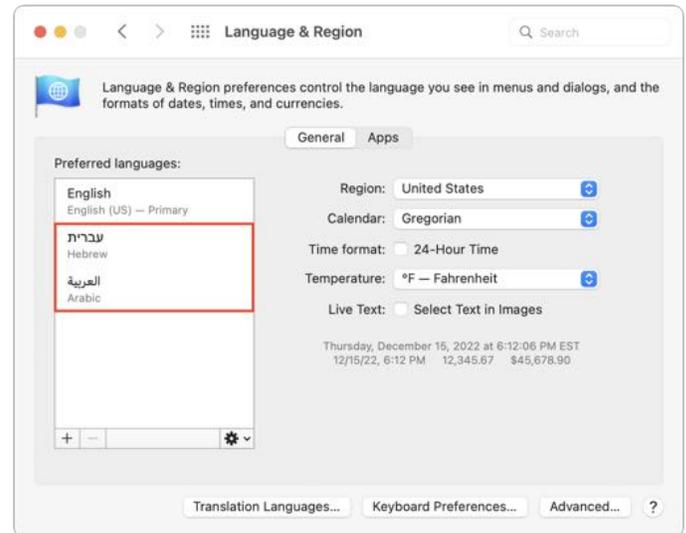
It turns out that Contacts on the Mac encloses all phone numbers in two zero-width (and thus invisible) Unicode characters: U+202D (LEFT-TO-RIGHT OVERRIDE) and U+202C (POP DIRECTIONAL FORMATTING). Dave speculated that those characters exempt the phone number from the surrounding text flow direction, which would be right-to-left with Hebrew and Arabic, for instance.

Contacts appears to store those characters with the phone number, but since they're zero-width, there's no way to avoid selecting them that I was able to find. I don't believe Contacts is adding those characters to the numbers as part of the copy because if you select a few characters in the middle of a number and then copy and paste them into Terminal, the invisible characters won't be present. They're also unique to the phone number fields; a number stored in the Notes field won't have them either. Finally, the issue is limited to Contacts on the Mac; the problem doesn't afflict phone numbers copied from an iPhone or iPad in my testing. (Some subsequent research turned up an Ask Different thread that describes the problem—which has been reported to Apple as a bug—and confirms that it still exists in macOS 13.1 Ventura.)

Why don't these characters cause more problems? When developers design apps that accept data, they generally create *input filters* that filter out unwanted data. A phone number field has no reason to accept invisible punctuation characters, so it's good programming practice to drop them. If nothing else, filtering user input is an essential security practice because attackers will often try to feed a system data to inject malicious code or to cause buffer overflows that could be exploited.

So while it seems weird that Apple designed Contacts to enclose every phone number in those

zero-width characters, it does protect phone numbers pasted into right-to-left text flows from being changed. It doesn't seem necessary for Contacts to do this for all Mac users, however. Apple could try to identify when it would make sense—perhaps only when the user has enabled a right-to-left language in macOS's Language & Region settings.

Practically speaking, if you run into a situation where a phone number copied and pasted from Contacts doesn't work as expected, the simplest fix is to type the number by hand. You could also paste it into Terminal and copy just the phone number again, avoiding the extra Unicode characters on either side.

If you regularly ran into this problem and wanted to rack up some geek points, you could create a Keyboard Maestro macro that automatically stripped those extra characters. Peeved by the need to retype numbers like an animal, Dave Fultz did exactly that, so you could use his macro as a starting point for whatever workflow makes sense to you.

Or you could do what I do and rely instead on the contact management app Cardhop, which doesn't suffer from the problem and is far more pleasant to use (see "Cardhop Puts Contacts Front and Center," 18 October 2017, and "Cardhop 2.0 Bundled with Fantastical in Flexibits Premium," 27 May 2021).

By Adam Engst

An Annotated Field Guide
to Identifying Phish

Do you like phish? Not the band, not tasty seafood dishes, and not the pretty tropical variety. I refer instead to the intellectual challenge of identifying phishing emails that attempt to get you to reveal personal information, often including login credentials or financial details, or entice you to call a phone number where trained operators will attempt to separate you from your money.

Phishing is a big deal, with a State of Phishing report from security firm SlashNext claiming that there were more than 255 million phishing attacks in 2022, a 61% increase from the year before. The Verizon Data Breach Investigations Report for 2022 says that only 2.9% of employees click through from phishing emails, but with billions of email addresses available to target, the raw numbers are still high.

Even before the LastPass breach (see "LastPass Shares Details of Security Breach," 24 December 2022) and the news about a data leak containing email addresses of over 200 million Twitter users, I had been noticing more phishing attempts evading Gmail's generally effective spam filters. I don't entirely blame Gmail for this—in many cases, I can see how the messages would be hard to catch.
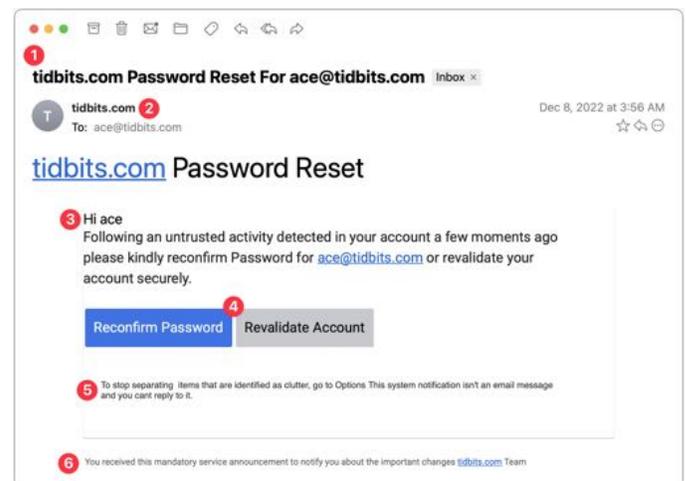
In the past, many phishing attempts were obviously fake, and intentionally so. That's because they only had to sucker people who were sufficiently inexperienced, credulous, or easily deceived that they would continue to go along with the scam. Now, however, I'm seeing phishing attempts that are more sophisticated and harder to identify quickly.

I've been examining phishing attempts for so long that it's hard for me to imagine what might fool someone else, so I wanted to share some recent attempts that slipped past Gmail's filters. For each

message, I've called out some of the ways I identified it as phishing. I suspect that most of you will assume that you would also easily have identified the message as fake, but remember, many people move rapidly through their email without reading carefully. Perhaps my calling out of some of the hallmarks of phishing attempts can help you or the people to whom you forward this article avoid being drawn in.

Password Reset

This phishing attempt purports to come from a system administrator in charge of my email domain and tries to lure me into clicking a button. The text isn't very good, but the buttons are, and it's easy to imagine someone who's scanning too quickly clicking the button without even reading the text. But that would be a mistake!

1. The Subject line for this message is glaring to me because I know how tidbits.com is managed—I do it! So this one wasn't going to fool me, but I could imagine someone getting a similar message that identified their domain and thinking it was from the IT department. Putting the email address in

the Subject might get some people to click because it feels personalized.

2. It feels weird to me that the From line of the message is also tidbits.com. Even if an administrator had to email users about a required password reset, I would expect a name or a role to appear here.

3. The body of the message is a giveaway. The phisher clearly doesn't know my name and is thus addressing it to my username. With additional data being leaked all the time, however, there's no guarantee that phishers won't start personalizing messages more thoroughly. The other problem with this text is that it doesn't sound like it was written by a native speaker of English. There's a difference between someone who's just a weak writer and someone who doesn't think in English, and the phrase "please kindly reconfirm Password" screams "non-native speaker." (I presume that phishing attempts are localized into other languages for people in other countries, so replace "English" with whatever your native language is.)
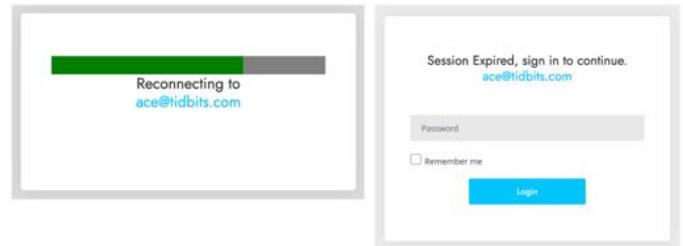
4. These buttons are pretty good—they look official, they're consistently capitalized, and one is clearly the default. But anything that contains the word "password" and tries to get you to click should set off warning bells.

5. The phisher should have stopped with the buttons. I adore this small text, which tries to claim that what you're looking at is a system notification, not an email message, and you can't reply to it. Yeah, sure, buddy.

6. Finally, we come to an even smaller line of text that sounds like an AI wrote it, in that it makes a nominally true but largely meaningless statement. I guess I could imagine a beleaguered system administrator writing something similarly banal, but to me, it grates.
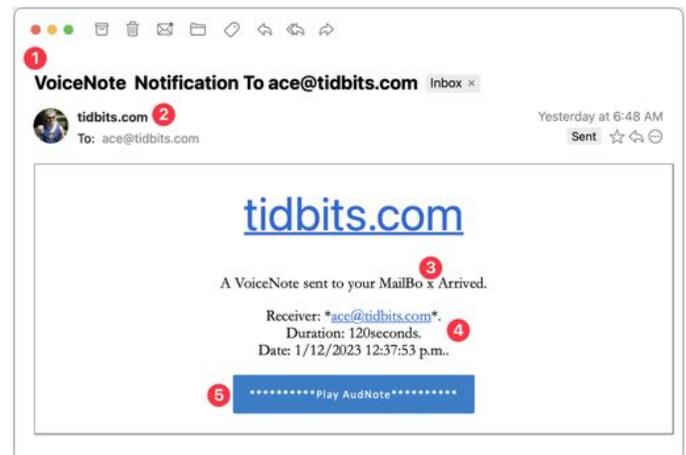
Interestingly, when I clicked either of those buttons (the risks I take for you, loyal readers!), I was directed to a Web page that showed the progress dialog on the left and then tried to get me to log in with the dialog on the right. I didn't take it any

further. It's worth emphasizing that you shouldn't click buttons or links in a phishing email since that might put your email address on a list of "dupes who click stuff" and invite even more phishing messages.

VoiceNote Notification

My next example pretends to be a notification of a voice note. Again, since I'm the person who would have set up any voice note system, this phishing attempt is wasted on me, but in a large organization where people don't know what IT is up to? Not inconceivable.
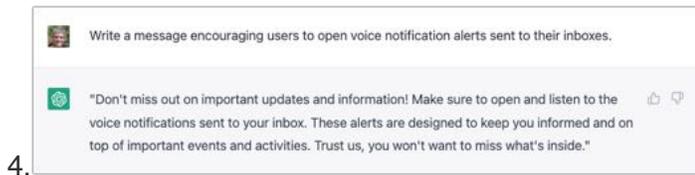

1. I could see someone opening this message because the Subject isn't problematic. I tend to think that if your email address ends up in the Subject line, it's probably spam, but I don't know if everyone thinks that way. The copyeditor in me also takes note of the two spaces between "VoiceNote" and "Notification" and the fact that "To" is capitalized. It's unprofessional and looks wrong, but real emails could have similar mistakes.

2. I wonder if this is the same phisher as in the previous attempt since it's again clumsily

attempting to fool me into thinking the message is from a system administrator at my domain. In fact, what's happened here is that the phisher has forged my address, which my email client, Mimestream, indicates by showing my avatar and putting a Sent tag over on the right. This message appears to be from me and to me, but many people wouldn't pick up on those subtle cues.

3. I'm not perturbed by the camel case in "VoiceNote," which matches the Subject and could identify some product or service. But the capitalization and extra space in "MailBo x Arrived" is glaringly bad. Had it said, "You have received a VoiceNote." more people might have believed it. I can't believe I'm making editing suggestions for spam, but with services like ChatGPT, phishers won't have trouble writing better text. They will have to be a little sneaky—ChatGPT tries to avoid helping miscreants.

4.

5. What's with the asterisks, the periods, and the lack of a space in "120seconds"? I'd expect cleaner text in a notification purporting to be associated with some VoiceNote service. The timestamp is also suspiciously precise—I can't see any real product designer allowing such detail.
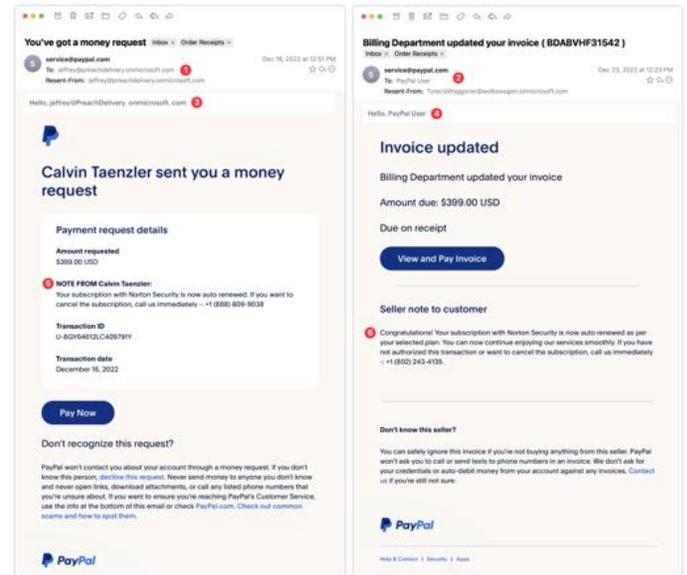
6. I'm so disappointed with this button! This phisher is far too enamored of asterisks, and they couldn't even stay consistent with the product name, switching suddenly to AudNote.

And in fact, when I click the lame button… drumroll, please…it is the same phisher as the previous one! Same progress dialog, followed by

the same login dialog. The first one arrived over a month ago, so it's interesting that the phisher is continuing to probe me. I wonder if clicking the links will cause more to come.

PayPal Money Request and Invoice

My next two examples are devilishly good because they use real PayPal messages, complete with multiple bits of text warning about phishing. But there are still giveaways. These arrived a week apart and are almost certainly the work of the same phisher.

1. The To line in the left-hand message identifies some random person. This is a glaring but odd mistake since phishers can generally get your address in the To line. The Resent-From line is also a clue because Mimestream is calling out the fact that the message was sent to and from the same person.

2. In the right-hand message, the To line merely calls out "PayPal User" and uses a different Resent-From line. That's less of a giveaway, but you should always be suspicious of messages that don't seem to know who you are or that are sent by people you don't know. (And if you know Tyrecia Waggoner, you have a tougher job.)

3. Another poor job of personalization on the left-hand message. Although the formatting of this line looks wrong to me, when I compared it
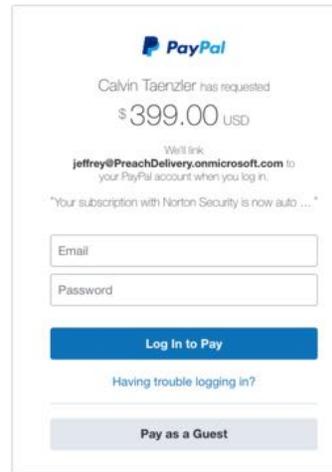
against a legitimate PayPal money request, it's the same, albeit with my name. If the phisher had used my name rather than the email address of the person who showed up in the To and Resent-From lines, it would have been more convincing.

4. I believe these two messages are from the same phisher, who is learning how to modify the PayPal messages better, since this one is addressed to "PayPal User." That's not too persuasive, but it's easy to imagine it becoming more so.

5. Who is [Calvin Taenzler](#)? Unless you know, you probably won't fall for this request, even after reading the fake message about your Norton Security subscription auto-renewing. I tried the phone number, but it wouldn't connect—perhaps it was live back in December when I received this phishing attempt.

6. The evolution of the PayPal scam continues. Instead of naming someone you won't know, the message pretends to be an updated invoice for that fake Norton Security subscription, and if you don't want it to auto-renew, you have to call a phone number. It's easy to imagine someone with failing memory thinking they might have subscribed to Norton Security and calling to check. Like the other number, it wasn't active anymore.

When I clicked the Pay Now button on the left-hand message, I was sent to a legitimate PayPal login screen. Needless to say, I didn't go any farther, but unless PayPal has shut this account down, it seems like it could be collecting money. Clicking the View and Pay Invoice button on the right-hand message also went to a legitimate PayPal page, but it displayed an error about the invoice being no longer available.
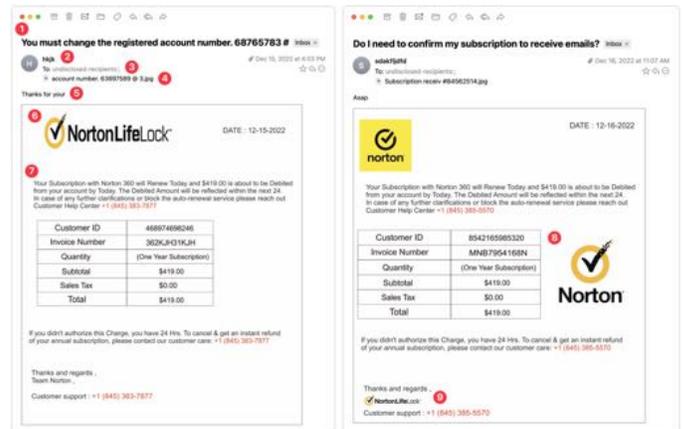

If you receive PayPal phishing attempts, [PayPal asks that you report them](#). Forward the messages to phishing@paypal.com and then delete them. *Do not mark them as spam*, or you might find legitimate PayPal messages being caught by your spam filter.

Norton 360 Invoice

The next two examples came a day apart and are very similar. Again, they seem to show a particular spammer—perhaps even the same one trying the Norton Security scam above—testing different techniques.

1. The Subject lines of these messages are both fairly effective—in both cases, I'd open them to see what they're talking about. Notably, neither has any egregious spelling, capitalization, or grammatical errors.

2. This phisher was lazy and just put some random letters in the From line. That's a dead giveaway, and had I been working normally, I would have marked these messages as spam instantly.

3. Notice that the To line says "undisclosed-recipients." On its own, that doesn't bother me much because it's not uncommon to receive Bcc'd messages that end up with something like this in the visible To line. However, these messages purport to be individual billing reminders, complete with customer ID numbers, and it doesn't make sense that a business that knows my ID number wouldn't also know my name.

4. Pay attention to the attachments here—both are JPEG images with names that don't set off any major warning bells. The invoices appear in the message windows because Mimestream displays attached images that way.

5. Without Mimestream showing the attachments in the body of the message, all you'd see is this one line of text: "Thanks for your" on the left and "Asap" on the right. That's how these phishing attacks evade Gmail's filters—there's too little to go on to mark the messages as spam automatically. We've all sent brief messages that are a line of text and an attached image.

6. On the left side, the NortonLifeLock logo is horizontally compressed. That's a tell for me—no large company with basic competence would allow email to go out with the corporate logo squished. It's different—and better—on the right.

7. The text is the same on both sides, apart from the phone number, and it's pretty bad. What is it about the intersection of English and whatever language so many phishers speak that they capitalize words with such abandon? If it looks like a message was suffering from a spasming Shift key, it's almost certainly not legitimate.
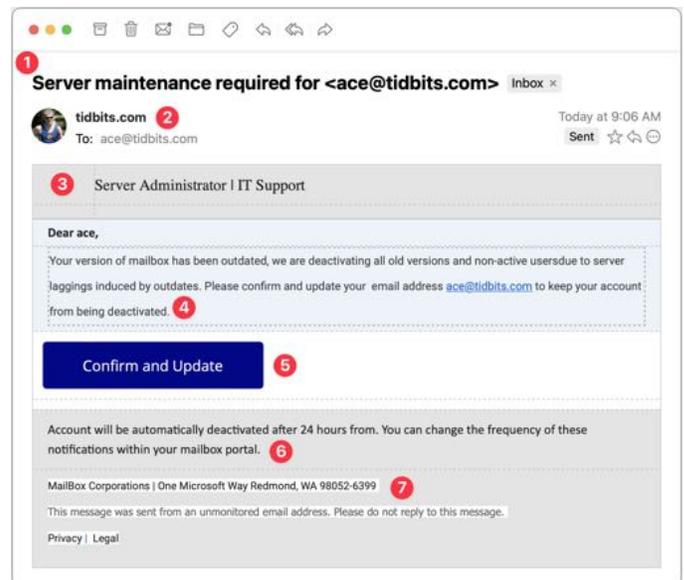
8. The table is a clever touch that makes the message seem more legitimate. I particularly like the row indicating that no sales tax was charged. Pay closer attention when a message starts to look real—a quick glance at a table like this could get you to click something if you weren't reading

carefully. Also, note that a day after the first message on the left, the phisher figured out how to avoid squishing the Norton logo and found a nicer place for it.

9. Lastly, you can see on the right side that the phisher has used the NortonLifeLock logo again, but as a signature of sorts, and with the proper aspect ratio. That looks much more real than the Team Norton text signature on the left. In short, beware of graphics. They're easily copied and are an easy way to make a message look more real.

Server Maintenance

For my final example, consider this message purporting to come from my company's IT Support department. It's trying to steal your email password because once an attacker controls your email, they can reset your password at other sites and compromise more sensitive accounts. Protect that email password above all else!

1. The Subject line in this message is once again pretty good. I'd open this message because even though I run all the TidBITS servers, they're all hosted at companies like Arcustech and DigitalOcean. I already receive notifications from those companies about planned outages and the like, so this Subject line doesn't set off many warning bells, apart from including my email address.

2. Since the message purports to be from my IT department, having the From line be the domain isn't seriously problematic, and the To line below is also correct. However, note the use of my avatar and the Sent tag off to the right. That's Mimestream indicating that the phisher forged the message to look like it came from me as well as being sent to me. Not all email programs provide such an indication.

3. The table formatting in this message—particularly the dashed lines—is rather odd. To me, the dashed lines scream "Fake!" although I don't know if others would pick up on that. Pay attention to gut feelings that something isn't right.

4. Once again, the text is poor, with numerous mistakes that indicate that the message was composed by someone who doesn't speak English natively or even bother to let auto-correct work. The claim that the IT department is deactivating accounts is designed to make you worry, but good IT departments wouldn't just send email to see if an account was in use before deactivating it.

5. This Confirm and Update button looks good and is an attractive target for someone scanning without reading carefully. Beware if you find yourself doing that!

6. This second bit of text is nonsensical—I would hope that anyone reading it would get that, but perhaps that's overly optimistic.

7. I'm amused that the phisher invented "MailBox Corporations" and located it at One Microsoft Way. It's a stylistic flourish but counterproductive if the message is supposed to emanate from my company's IT support department. Mismatches like this are another good hint that a message isn't real.

Clicking the Confirm and Update button loaded a generic cPanel login for a webmail service. It has a little trouble with its copyright symbol, and I hope it wouldn't fool anyone who was expecting to see their organization's webmail system. I fed it a random email address and password and got an expected error, but I presume that all the data is being captured on the other side.

Overall Advice

I hope you've found this tour of a few selected phishing attempts helpful. As you can see, they vary widely in their sophistication and techniques, but with careful attention, you shouldn't have much trouble resisting their siren calls. My main pieces of advice for identifying phishing attempts are:

• Assume that you shouldn't click anything in an email message until you've given it a closer look. It's easy to skim and respond to a full inbox, which is exactly what phishers are counting on.

• Read the text of the messages, looking for capitalization, spelling, and grammatical mistakes. Nothing prevents phishers from writing correct English, but it seldom happens. Pay special attention to the fine print at the bottom—it's often more of a giveaway than anything else because phishers are trying to make the message look right without continuing the deception.

• Evaluate any claim about something happening within your organization against what you know to be true. Would your IT department do something extreme like deactivating accounts with little warning? Even if that's not inconceivable, it's safer to ask someone in the organization if a message is real than to click buttons randomly.

• Be careful when you see big, legitimate-looking buttons. They're easy to make and can sucker people who don't read the surrounding text carefully.

• If a message is just an attached image, it's almost certainly fake.

• With messages that don't set off any other warning bells, like the PayPal phishing attempts above, be alert for names and email addresses that aren't familiar.

Good luck, and stay safe out there!