

## **Google is obsessed with phishing, thankfully**

People associate **phishing** with identity theft. That's bad enough, but there is something else to consider. If phishing continues to be successful, people will be afraid to do anything online, especially when it requires disclosing personal information. Businesses, financial establishments, and companies who exist because of the internet are keenly aware of this. Google has decided to bring their vast arsenal of technology to bear on the problem of phishing.

There are two reasons to be interested in Google's approach to anti-phishing. First, their anti-phishing team has been able to automate the black-listing process, no small feat. Second, they are finally talking about how they do it. Their method involves two parts, a client-server interface and the backend data base. Let's look at the client-side service first.

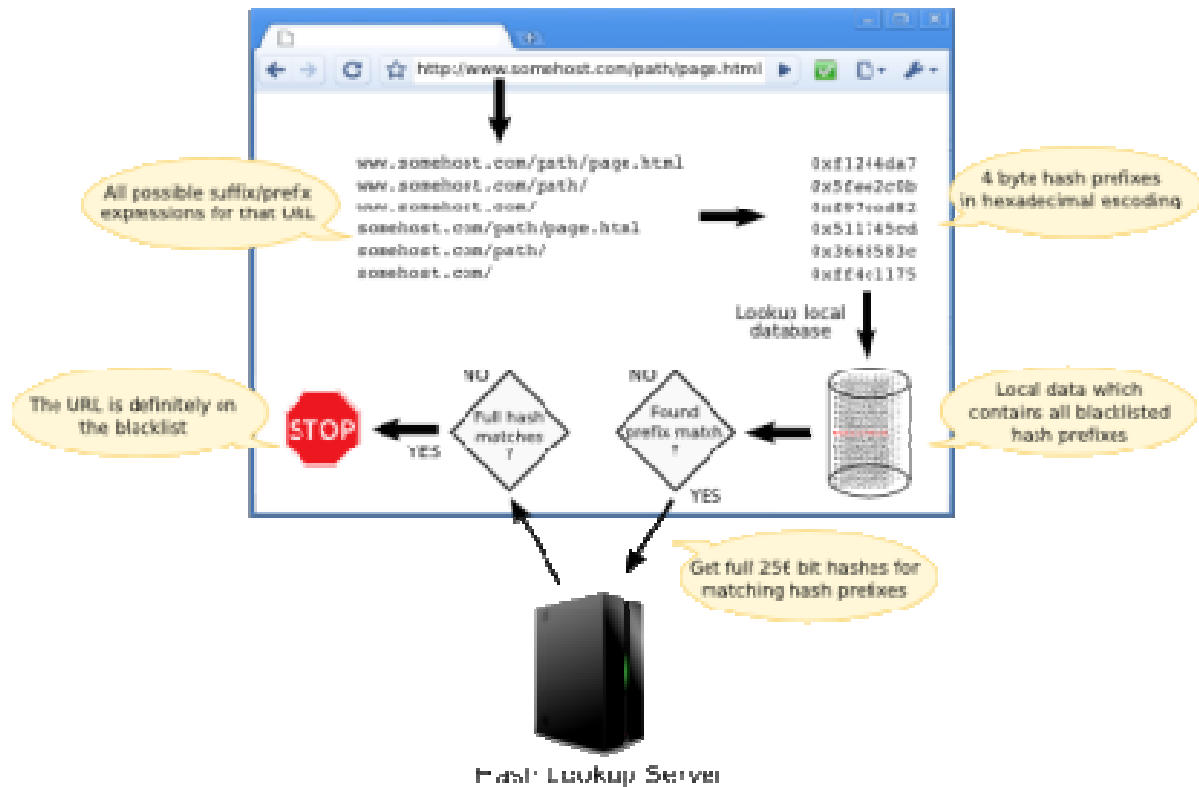
### ***Google's Client side Solution***

The client interface is Google's **safe browsing api**. Many do not realize that it has been working quietly in the background for several years. Not only that, but three well-known web browsers: firefox, safari, and chrome, use it. Google defines the safe browsing service as:

*“at a high level, the service works by checking each url the client loads against a list of known phishing and malware sites. The list of known sites is represented as host-suffix / path-prefix expressions.*

*As the name suggests, these expressions can match arbitrary urls as long as they have the required host suffix and path prefix. This approach helps protect against sites where the attacker uses many different urls in order to try to evade blacklists.”*

The following diagram (courtesy of google) is a visual description of the look-up process:



The client-server handshake is the easy part. Trying to keep the black list current, have minimal mistakes, and even fewer false positives is where it gets tricky.

### ***Google's Back End Solution***

Until recently, Google has kept mum about how their black list is populated. Three members of Google's anti-phishing team: Colin Whittaker, Brian Ryner, and Marria Nazif published a paper describing their methods at the **17<sup>th</sup> annual network and distributed system security symposium**. Two links describing anti-phishing approaches are:

Google's approach: <http://tjscott.net/446securityplus/google.phishing.paper.pdf>

The team discusses what is needed for the black list to be effective:

- **Comprehensive:** a blacklist that is not comprehensive fails to protect a portion of its users.
- **Error free:** false positives subject users to unnecessary warnings. Eventually, the users will ignore the warnings.
- **Timely:** the black list must update in real-time. As most phishing sites are up for less than a day.

They go on to explain that the automatic classifier or “back-end algorithmic process” uses the following web-page elements in the decision-making process:

- **Page url:** look for anything odd about the hostname. Is it unusually long or possibly contain an ip address.
- **Page content:** the page is checked to see if it has a password and or pin field. Additionally, the page is checked for links that may be pointing at a known phishing domain.
- **Tf-idf score:** tf-idf is a ranking method used when automatically scanning for phishing sites. The mathematics behind this gives more weight to important terms like “password” or “pin”.
- **Hosting information:** what network hosts the web site and where the web servers are located geographically can be telling. For example, what if a web server for an American bank is located in a different country.
- **Pagerank:** pagerank is used to determine the spam reputation of the page’s domain. Apparently, the anti-phishing team has discovered a relationship between phishing pages and domains that send spam.

As you can see, Google checks a lot of things.

### ***Automated Processes***

Automating the search for the above elements was the first and probably the simplest step the team did. The classifier then takes the information and ranks the url, from 0.0 not at all phishing to 1.0 definitely phishing. Finally, software called the “**blacklist aggregator**” prepares the list to be served to the clients.

What really makes this system effective is how the classifier is retrained every day to pick up new phishing trends. Google explains:

*“As a training data set, we use a sample of roughly ten million urls analyzed by the classification workflow over the past three months along with the features obtained at the time.”*

The report goes on to explain how the training data set is manipulated to test the classifier and make sure it is providing the most accurate results possible. From what i understand, the training process is the heart of the classifier and what separates google’s approach from others.

### ***Another Google Concept***

Another **google blog post** explains how web-site designers can minimize the chance of having their work trigger anti-phishing scanners. The post explains that we should keep the following points in mind:

- Beware of username and password requests that are not specifically for that web site.
- Be leery of logos near login fields that are not related to the web site.
- Links to other web pages should be readily viewable and related to the site’s domain page.

The above bullet points are important, but easily missed. I was almost tricked by a password request that had nothing to do with the web site i was viewing.

### ***Some Conclusions from the Google Report***

There are so many anti-phishing rules that it is difficult to keep track of all of them. Users often reject **security advice**, such as how to implement anti-phishing capabilities for they are often just too complicated. Google’s approach is heartening and their conclusions offer the following encouragement:

*“In this paper, we describe our large-scale system for automatically classifying phishing pages which maintains a false positive rate below 0.1%. Our classification system examines millions of potential phishing pages daily in a fraction of the time of a manual review process. By automatically updating our blacklist with our classifier, we minimize the amount of time that phishing pages can remain active before we protect our users from them.*

*Even with a perfect classifier and a robust system, we recognize that our blacklist approach keeps us perpetually a step behind the phishers. We can only identify a phishing page after it has been published and visible to internet users for some time. However, we believe that if we can provide a blacklist complete enough and quickly enough, we can force phishers to operate at a loss and abandon this type of internet crime.”*

### ***Final thoughts***

Automated filters aimed at reducing phishing attacks are vital to the existence of the internet as we know it. There may be other answers, but until they are turned into working systems, this seems like our best bet. When a security topic is so important, researchers work towards implementable, and then, robust solutions. A user-oriented solution called ItrustPage is shown below:

<http://tjscott.net/446securityplus/itrustpage.pdf>

Famous advice from Sun Tse, in the Art of War, can be paraphrased as “if you know both yourself and your enemy, in 100 battles, you will not lose one.” Of course, we translate “yourself” as our network, and “your enemy” as phishers. It should be clear that the more informed we are about phishing, the safer we will be. Using firewall language, more knowledge about phishing would be a “belt and suspenders” approach.