On Saturday 10 April 2020, the Open Tree of Life web servers were attacked by hackers. We will use this page to post information about the attack as we learn more.

What was the timeline of the attack? The first evidence of unusual activity that we have found was at 3:15PM UTC. By 6:00PM UTC the attack had degraded system performance enough to cause our automated tests to fail. We had a hard time connecting to the servers, but as soon as we got on them, we realized there had been a security breach. So we took all of the servers offline at around 8:30PM UTC.

What were the attackers doing? Several of our servers were attacked. The hackers were able to execute code on the servers that was being used to perpetrate a denial of service attack on another web site (the target was not related to Open Tree or [as far as we can tell] in any way related to bioinformatics). A denial of service attack consists of hitting a web server with many simultaneous requests so that it can no longer function correctly. Because it can be expensive and illegal to run such an attack, hackers frequently infiltrate other systems (the Open Tree of Life servers, in this case) to launch their attack. As far as we can tell, the hackers were not interested in the Open Tree project in any way, they were simply using our bandwidth and CPU resources to attack someone else. All of the affected servers appeared to be running the same script to attack the same third party website.

How did the attackers get in? We are still diagnosing the specifics, but we were using an older version of one web framework. It appears that there are some vulnerabilities in that version, and this appears to be the most likely route that the attackers used. All of our servers that were using that framework were breached, and none of the servers that were not using that framework were breached (though there is evidence that attackers tried to get into these, as well). For that reason, we think that the hackers noticed a vulnerability on one machine, then attacked all of the servers in the opentreeoflife.org domain, and were successful on the servers that were running the old framework.

Are there risks to Open Tree users? By design, Open Tree collects very little information from users and we don't maintain private databases of users information. We don't require registration or authentication to visit the site or call its APIs. Curators and commenters authenticate through a GitHub API that gives opentree a token that allows the server to make commits to public GitHub repositories in the name of the curator. This is how we give credit to curators as they input trees in the phylesystem repository (https://github.com/OpenTreeOfLife/phylesystem-1/commits/master). Because the attacker could launch a process on the servers, it is possible that they could have gained access to some of those tokens if they were very aware of details of the architecture of the Open Tree project. At around 10:00PM UTC on Saturday (after we shut down the servers) we revoked all of these user tokens. We do not think that the attackers were after Open Tree user data; all of our investigations point to the servers just being used for this generic denial of service attack. However, if you were an Open Tree of Life curator and notice any GitHub commits to your public repositories that you did not make on Saturday, 10 April, 2020, please let us know. You can place comments that you do not mind being publicly visible at https://github.com/OpenTreeOfLife/feedback/issues. If you have issues that you would like to communicate in private (as private as email ever is), feel free to email me at: mtholder@ku.edu

What is Open Tree of Life's plan to get online and avoid future hacks? We are working hard on several steps:

1. We are hoping to restore access to the public APIs that do not use the vulnerable web framework. This will restore the tree-of-life, taxonomy and tnrs portions of our web services. ✅

2. The front-end of the web site and the portion of the API involved in curation of studies will be upgraded to the latest version of the web framework. ✅
3. We will tighten our access rules to prohibit access to URLs that are not used in standard Open Tree of Life web calls. ✅
4. We are installing another layer of security software on the servers to help us detect unusual activity, and we will increase the frequency of our monitoring efforts. (partial ✅ we are monitoring more thoroughly, and have installed some of the open source security software. We are still looking at augmenting that set of software).
5. The steps mentioned above will allow us to start serving a read-only version of the site. ✅
6. We are migrating to a new system for working with our phylesystem repository of studies that will only grant a much restrictive set of permissions to the Open Tree servers. This work has actually been underway for a little while. We will not reenable curation or commenting on the web servers until that newer set of user permissions is implemented. ✅

~~At this point we do not have good time estimates for these deadlines, but we will update this page as we have more confidence in those timing.~~

Update in early June 2020: we have re-enabled curation using new less expansive permissions required from curators.

We sincerely apologize for the down-time and the security breach,
Mark Holder on behalf of the Open Tree of Life team