

From: [REDACTED] <[REDACTED]@MIT.EDU>
Sent: Wednesday, October 20, 2010 2:01 PM
To: [REDACTED] <[REDACTED]@ithaka.org>
Cc: [REDACTED] <[REDACTED]@mit.edu>; [REDACTED] <[REDACTED]@mit.edu>;
[REDACTED] <[REDACTED]@mit.edu>; [REDACTED]
<[REDACTED]@ithaka.org>
Subject: RE: 10:00 am Update: JSTOR Abuse at MIT: All IPs Blocked

Hello [REDACTED], I apologize for the delay in responding. I have been waiting for more detail from our IS&T group, but I still don't have that information. In the meanwhile, I am hoping we can move forward with our discussions to make the technical changes necessary to implement our additional authorization layer, which we call "econtrol."

We would like to try an additional tweak to our normal 'econtrol' process if it is possible on the JSTOR end. If JSTOR could use an Apache mod_rewrite to redirect any activity from the MIT IP addresses (aside from those for our proxy server) to our proxy server, our patrons would not have to remember to use a special gateway to get to JSTOR. This would be a big benefit to our patrons. Would this be feasible for JSTOR?

When we are ready for the change, we will need you to reduce the authorized MIT IP ranges list you appended below to these:

18.51.1.222
18.7.29.240

Ultimately only the second address will be needed, but we are in transition from one proxy server to another. If you can implement the server-side configuration change mentioned above, we would currently be redirecting to 18.51.1.222, though we'd have to change to the other address in the next few months.

Please do not make any change yet, as we want to send information to our patrons before the switch so that they will be aware of the changes in the access model. We are preparing the communications now, but we need to know whether JSTOR can implement the apache mod_rewrite before we finalize those messages, since that information will determine what we have to tell our patrons about access.

I have copied [REDACTED] from our IS&T group here -- if there are technical questions about the Apache mod_rewrite you will be best served by going direct to [REDACTED]. I have also copied the [REDACTED], [REDACTED], as [REDACTED] will be the one making the changes on our end to EZproxy. Those changes will need to precede the changes on JSTOR's end.

In at least partial answer to your inquiry below, we offer guests access to the MIT network. However, once we institute our additional authorization layer for JSTOR, this route will be closed to guests. So we will have closed the pathway through which the excessive use occurred.

We look forward to moving forward with an econtrol implementation with or without the apache code --

Thanks very much,
[REDACTED]

██████████
████████████████████
MIT Libraries

P ██████████
██████████@mit.edu

<http://libraries.mit.edu/scholarly>

-----Original Message-----

From: ██████████ [mailto:██████████@ithaka.org]
Sent: Monday, October 18, 2010 10:04 AM
To: ██
Cc: ██████████
Subject: RE: 10:00 am Update: JSTOR Abuse at MIT: All IPs Blocked

Thanks ██████████,

I appreciate your candor here. I have dealt with many cases over the years and understand the difficulties inherent with tracking down individual users. I am hopeful we can use this opportunity to work together towards building more stable, sustainable and secured access to JSTOR. We are meeting as a larger group to discuss this matter further this afternoon and I am wondering if you could provide one point of clarification for that conversation.

Understanding that you may not be able to identify the individual, where you able to identify the credentials used to access MIT authorization for this activity? That is, was there a shared UN/PW used for guests or an open port on a proxy used in this case. Basically, the concern is, as we sort out the IP configurations necessary with you, could this or any user use the same authorization methodology to do this again or has the pathway been identified and locked down?

As for your IP configurations and establishing an access point within your range. We have the following IP addresses currently installed for MIT. Please let me know which ones to maintain and which to remove as needed and we'll get right to it.

- 18.*.*.*
- 128.30-31.*.*
- 128.52.*.*
- 129.55.*.*
- 192.52.61-66.*
- 198.125.160-163.*
- 198.125.176-192.*

Best,

██████████
████████████████████

JSTOR | Portico

rooted in our deep investment in our ability to be good stewards of the content in our care. Noting that, the 9/25 and 9/26 incidents yielded more than 400K PDFs from JSTOR to this user. This is both unprecedented and very concerning. Any progress you could provide on identifying the user responsible and the steps I mentioned yesterday would be very helpful.

We are also mindful that the weekend is upon us and are looking for collaboration on your end and acknowledgement that MIT staff will be monitoring your systems closely for any recurrence, as we are, until we can reach resolution.

Incident on 9/25 & 9/26

IP = 18.55.6.215
Start = 25-SEP-10 05:06:49.109524 PM
End = 26-SEP-10 04:24:54.297995 AM
Total Sessions = 1,256,249
Total Articles Downloaded = 453,570
Total Journals Affected = 562

Incident on 10/9

IP = 018.055.005.100
Start = 2010-10-09 14:53:18 from
End = 2010-10-09 19:08:01
Total Sessions = 8,515
Total Articles Downloaded = 8,422
Total Journal Affected = 714

Best,

██████████

-----Original Message-----

From: ██████████ [mailto:██████████@MIT.EDU]
Sent: Thursday, October 14, 2010 12:44 PM
To: ██████████
Cc: ██████████
Subject: RE: 10:00 am Update: JSTOR Abuse at MIT: All IPs Blocked

Hello ██████████ and ██████████,

Our investigations here point to the same guest that was involved in the 9/27 incident. We don't have enough information to follow the trail completely, but the signs suggest that the same guest user was responsible for this

latest activity. To pursue this further, our IS&T group would need more information. Specifically, they are wondering if you are seeing any robotic activity from MIT currently and if so, whether you have any information about the IP addresses involved.

Given that it appears all of this excessive use was caused by a guest visitor at MIT, we have been considering next steps, and would like to suggest that we move to a new access model that will eliminate use by guests. We have recently developed an additional authorization layer that we can apply to particular products to prevent access by guests/walkins. We've tried this approach with one or two publishers where we had seen repeated excessive use, and it has stemmed the problem in those cases.

We would orchestrate this change by changing the proxy configuration on this end, and then we'd ask you to change the list of acceptable MIT IPs to only our proxy server's address -- a single IP.

If this sounds like an acceptable approach, let's discuss the next steps. To carry out the change, I'd have JSTOR work with [REDACTED], copied here.

Best,

[REDACTED]

[REDACTED]
[REDACTED]
MIT Libraries

P [REDACTED]
[REDACTED]@mit.edu

<http://libraries.mit.edu/scholarly>

-----Original Message-----

From: [REDACTED] [mailto:[REDACTED]@ithaka.org]

Sent: Tuesday, October 12, 2010 10:09 AM

To: [REDACTED]

Cc: [REDACTED]

Subject: 10:00 am Update: JSTOR Abuse at MIT: All IPs Blocked

Hello Again,

We have requested that the IP range be unblocked at the firewall and that process is currently underway. I will confirm when that is accomplished and report the IPs and timestamps surrounding the event shortly.

[REDACTED]