ON THE JOINT ANALYSIS REVIEW, AKA THE FALSE TOR NODE POSITIVES REPORT

As I noted here, everyone agrees that the Joint Analysis Report released with Obama's sanctions package is a shitshow (here's the best explanation of why). But aside from complaining about how the shitshow JAR undermines the Administration's claims to have confirmed Russia's role in the DNC hack, no one has tried to explain why the Administration would release such a shitshow report.

Until now. Jonathan Zdziarski argues that the reason the Administration released a shitshow report is because they're very worried about the extent of Russian infiltration in our infrastructure, and by releasing a bunch of indicators that a probably not Russians but might be, it will get a lot of people (like utility Burlington Electric) looking for things that might be Russia, all while protecting the real intelligence that would expose sources and methods.

> One thing that's been made clear by recent statements by James Clapper and Admiral Rogers is that they don't know how deep inside American computing infrastructure Russia has been able to get a foothold. Rogers cited his biggest fear as the possibility of Russian interference by injection of false data into existing computer systems. Imagine the financial systems that drive the stock market, criminal databases, driver's license databases, and other infrastructure being subject to malicious records injection (or deletion) by a nation state. The FBI is clearly scared that Russia has penetrated more systems than we know

> about, and has put out pages of information to help admins go on the equivalent of a bug bounty.
>
> Everyone knows that when you open a bug bounty, you get a flood of false positives, but somewhere in that mess you also get some true positives; some real data. What the government has done in releasing the JAR is made an effort to expand their intelligence by having admins look for (and report) on activity that looks like / smells like the same kind of activity they found happening with the DNC. It's well understood this will include false positives; the Vermont power grid was a great example of this. False positives help them, too, because it helps to shore up the indicators they're using by providing more data points to correlate. So whether they get a thousand false positives, or a few true ones in there, all of the data they receive is helping to firm up their intelligence on Russia, including indicators of where Russia's interests lie.
>
> Given that we don't know how strong of a grasp Russia has on our systems, the JAR created a Where's Waldo puzzle for network admins to follow that highlights some of the looser indicators of compromise (IP addresses, PHP artifacts, and other weak data) that doesn't establish a link to Russia, but does make perfect sense for a network administrator to use to find evidence of a similar compromise. The indicators that tie Russia to the DNC hack were not included in the JAR and are undoubtedly classified.
>
> There are many good reasons one does not release your evidentiary artifacts to the public. For starters, tradecraft is easy to alter. The quickest way to get

> Russia to fall off our radars is to tell them exactly how we're tracking them, or what indicators we're using for attribution. It's also a great way to get other nation states to dress up their own tradecraft to mimic Russia to throw off our attributions of their activities. Secondly, it releases information about our [classified] collection and penetration capabilities. As much as Clapper would like to release evidence to the public, the government has to be very selective about what gets released, because it speaks to our capabilities. Both Clapper and Congress acknowledged that we have a "cyber presence" in several countries and that those points of presence are largely clandestine. In other words, we've secretly hacked the Russians, and probably many other countries, and releasing the evidence we have on Russia could burn those positions.

I don't know. I remember that Khalid Sheikh Mohammed had the CIA chasing black Muslim extremists planning to set forest fires in Montana for three months. False positives waste limited resources. Perhaps the intelligence community thinks this is okay because it's not *their* resources that will go to waste. But the entire thing seems to have increased the skepticism about the value of the government's threat reporting, which is all in all a bad thing.

But false positives do have two other purposes. I would hope these two aren't the reason why the IC released a shitshow report, but it deserves consideration.

First, false positives raise the fear level. Last week's Vermont false alarm is the perfect example of that: within hours — even on a Friday night — much of the country was worrying about our power grid. And remember, that false alarm was leaked by a Senior Administration Official

that chose to leak it to someone who is not an expert in this field.

At that level, this felt like the 2004 leaks about an election year al Qaeda plot that — we now know — were secretly used to reauthorize torture and the dragnet, but which were largely bogus and partly based off torture. I can only imagine the kind of heightened surveillance the IC is putting in place behind all this fearmongering.

But there's another effect of the false positives that have *already* been generated by this report: tying a bunch of Tor nodes to Russian spying. Almost immediately after the report came out, Jerry Gamblin found that 21% of the IP addresses were Tor nodes. Micah Lee did more analysis and found that 49% of the IP addresses in the report are or recently have been Tor nodes.

What we don't know about the Tor nodes, though, is how they came to be included in the report. Did they just happen to be used in a Russian attack; did the Russian hackers just let Tor randomly assign which node they exited from?

Or did the hackers choose — as you can do — which nodes they might use? There are a few reasons to pick a certain node over another. If you're trying to watch the Beeb's coverage of the Olympics, for example, you've got to pick a node in England.

But a more likely choice, for a smart Russian hacker, is to selectively choose nodes that the hacker believes would not keep logs.

Now consider some of the nodes that have been identified specifically. A Dutch paper made a big stink that the node operated by Rejo Zenger, who works at Europe's equivalent to EFF, was on the list. Something like 11 of the IP addresses are nodes operated by Calyx Institute, the non-profit ISP operated by Nick Merrill.

Merrill is, as you may remember, the guy who spent a decade challenging a National Security

Letter he received back in 2004. A big part of what he exposed is that the FBI was wrongly trying to get data flow with NSLs. In the last year, spooks have made several, thus far unsuccessful, efforts to get legal sanction for what Merrill exposed, the illegal acquisition of Electronic Communication Transaction Records using just an NSL.

Maybe Russian hackers chose to exit through Merrill's Tor nodes because he doesn't log traffic. Or maybe the government included him on this list because they know he doesn't log traffic.

The effect, however, is to (temporarily) burn select Tor nodes, perhaps those that don't log traffic, making it harder for anyone the government is trying to pursue through Tor to use it (and probably also making it more likely they'll use one of the many nodes believed to be operated by US intelligence). We know the NSA does a variety of things to force traffic onto switches it has access to; could the JAR just be a very elaborate way of forcing Russian traffic onto Tor nodes the FBI and NSA have access to?

Not to mention tarring the most committed privacy activists with association with Russian hackers.

Maybe that's not the intended effect of a report designed to generate false positives. But I'm sure the government considers it a happy side effect.

Update: Sounds like just about everyone found these indicators in their logs.

> Robert M. Lee, CEO of the Maryland-based industrial security firm Dragos Inc., warned his customers, who span critical infrastructure including water, electric, manufacturing and petro-chemical sites, that the technical information was bad. About one dozen called with concerns.
>
> "Every single company we have as a

> customer who ran the indicators got alerts, and all the alerts were bad," Lee said. "These addresses were not only not descriptive of Russian activity, they were not descriptive of malicious activity. They were actually common sites."
>
> [snip]
>
> One of the businesses that called Williams reported that an address tracked to Microsoft's telemetry server, which sends data to Microsoft when an application crashes. That conversation with his client spun into an hour-long discussion of "can we trust this report at all?" Williams said. "My short answer on this is no."
>
> He added: "This has a real cost to business. I suspect for a lot of them there (was) a lot of money spent chasing ghosts."