For Good Measure

Is the Cloud Less Secure than On-Prem?

DAN GEER AND WADE BAKER

Dan Geer is the CISO for In-Q-Tel and a security researcher with a quantitative bent. He has a long history with the USENIX Association, including officer positions, program committees, etc. dan@geer.org

Dr. Wade Baker is an Associate Professor in Virginia Tech's College of Business, teaching courses for the MBA and MS of IT programs. He's also a Co-Founder of the Cyentia Institute, which focuses on improving cybersecurity knowledge and practice through data-driven research. Prior to this, Wade held positions as the VP of Strategy at ThreatConnect and was the CTO of Security Solutions at Verizon, where he had the great privilege of leading Verizon's annual Data Breach Investigations Report (DBIR) for eight years. wbaker@vt.edu

> So you got to let me know,
> Should I stay or should I go?
> —The Clash

According to Deloitte's Chief Cloud Strategy Officer, "[2019] is the year when workloads on cloud-based systems surpass 25 percent, and when most enterprises are likely to hit the tipping point in terms of dealing with the resulting complexity" [1]. Given the nature of For Good Measure (this column), it may surprise you that it wasn't the 25 percent statistic that caught our attention in Deloitte's quote; it was reference to a "tipping point" where "dealing with the resulting complexity" in the cloud begins to negatively affect security. So we ask, do we see evidence that this is occurring? Are the rate of security exposures in the cloud higher than on-prem?

Conducting such an analysis requires data on security exposures affecting both on-prem and cloud-based hosts. RiskRecon [2] was kind enough to provide a sanitized data set derived from their efforts to provide visibility into third-party cybersecurity risk. For each organization analyzed, RiskRecon trains machine-learning algorithms to discover Internet-facing systems, domains, and networks. For every asset discovered, RiskRecon analyzes the publicly accessible content, code, and configurations to assess system security and the inherent risk value of the system based on attributes such as observable data types collected and transaction capabilities. The data set supplied by RiskRecon spans 18,000 organizations and over five million hosts yielding 32 million security findings of varying severity. Digging in, what can we determine about what organizations are seeing with respect to security complexities in the cloud vs. on-prem?

Figure 1 offers a bird's-eye view of our leading question. Each dot represents an organization in our data set, with a sufficient number of hosts in both on-prem and cloud environments to support this test. Their position on the grid is the intersection of the percentage of on-prem (horizontal) and cloud-based (vertical) hosts that have high or critical security findings. So, for example, the firm indicated by the arrow has an on-prem exposure rate of approximately 8% compared to a much lower 0.2% in the cloud. Organizations marked by blue dots (below the line) indicate they have comparatively fewer security issues when in the cloud. Green dots (above the line) represent firms that appear to be better off on-prem. Overall, there's a 60/40 split between organizations that operate with fewer issues on-prem (60%) vs. in the cloud (40%).

We infer from these results that the question of security destiny in the cloud is not predetermined. If you go, there may indeed be trouble; if you stay it may or may not be double. And it very well could be half.

Unfortunately, we do not have historical data available to determine whether those numbers are trending toward or away from a 50/50 "tipping point," but we were able to identify some

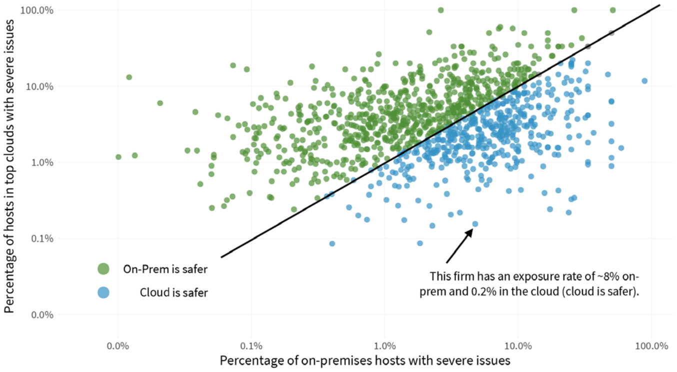

Figure 1: Comparison of hosts with severe findings in on-prem vs. cloud environments. Dots above the line indicate firms that have comparatively fewer security issues when on-premises.

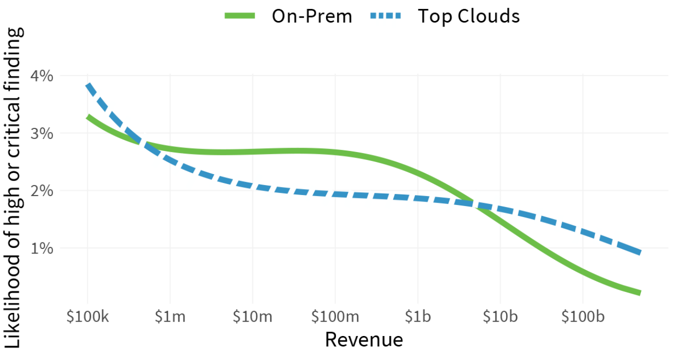

Figure 2: Models comparing exposure rates on-prem vs. cloud by organization size (annual revenue in log scale)

factors that affect a firm's likelihood of landing on one side of that line or the other. We discuss three of these factors below.

The Deloitte quote provides inspiration for the first factor we wanted to investigate. There's an implied statement that higher cloud adoption leads to a tipping point where added complexity affects security. Do we see evidence in the data that such a tipping point exists? To test that, we compared the rate of high and critical security findings in the cloud with the percentage of all hosts in the cloud for each organization. The result was a statistically significant but very low positive correlation (r=0.07) between those two variables. In other words, security exposures do increase as organizations put more and more hosts in the cloud…but not by much and only gradually. Not exactly evidence in favor of a tipping point.

The second factor is organization size as measured by annual revenue. We'd like to more directly measure characteristics like resources, IT complexity, and security capability, but size is the best proxy we have for those things. The question in view here is whether firm size (revenue) increases or decreases the likelihood of severe security exposures in cloud and on-prem hosts. Figure 2 constructs a regression model to test this correlation.

Let's first observe the general trend of decreasing likelihood of exposure as revenues grow for both on and off-prem hosts. This may reflect increased resources and maturity but may simply be an artifact of scale. It's almost inevitable that the likelihood of any single host being exposed declines as total population grows in larger enterprises.

Beyond that general trend, Figure 2 reveals some interesting "tipping points" between security in the cloud and on-prem. According to the model, organizations with annual revenues between $1M and ~$5B operate a little more safely in the cloud. The opposite holds true for firms outside that range—the really small and the really big. Might this imply that fast-growing organizations will want to use the cloud preferentially, but not small organizations and not giant, established players?

The third and final factor looks at the effect of consolidation vs. diversification in the cloud. In other words, is it better from a security perspective to consolidate hosts into one (or a small number of) cloud provider(s) or to spread services across many providers? Figure 3 reflects the data's answer to that question.

The "bars" in Figure 3 are actually made up of "dots" representing the 18,000 firms in our sample. We visualized it this way to emphasize the high degree of variation among organizations, especially toward the left side. But our focus is on the trendline, which turns out to be quite interesting. It suggests that the rate of severe findings is at its highest when cloud diversity is at its lowest. As organizations use more cloud providers, that rate drops steadily…to a certain point. Firms with four clouds exhibit one-quarter the exposure rate of those with just one cloud provider. Having eight clouds drops that rate in half again. Beyond that, security issues level off and even begin to rise among hyper-diversified cloud users. We can't help but see a kind of "tipping point" here: there's a point where consolidation and diversification find balance in the cloud, and that point varies from firm to firm. Echoing Deloitte, is that balance where complexity and the ability to manage it are themselves in balance?

One bit of caution regarding Figure 3: all kinds of factors are at play here that we cannot consider in our analysis. For instance, perhaps many of the firms with only one cloud provider are simply experimenting. This may reflect various stages of cloud maturity from left to right rather than the effects of consolidation vs. diversification. Given what we learned from Figure 2, one may hypothesize that this simply reflects the effects of organization size on exposure rates (the assumption being larger enterprises use more clouds). We included both variables in our analysis, but the number of cloud providers alone was the significant one.

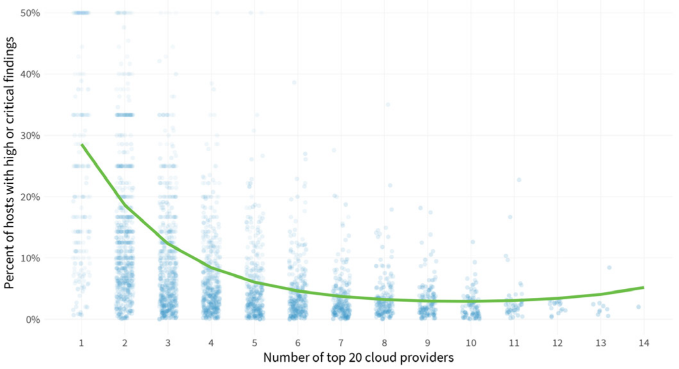

Figure 3: Rate of security exposures among hosts by number of cloud providers

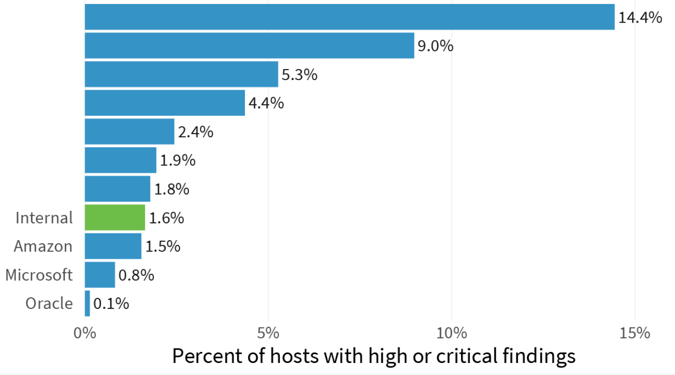

Figure 4: The spread in insecurity across major cloud providers

Of course, not all clouds are the same, either, as illustrated by Figure 4. Here we compare the prevalence of severe security findings among the top cloud providers. "Top" here refers to adoption. The clouds represented in Figure 4 accounted for over 90% of the cloud-based hosts in our data set. We also include the comparable rate for internal (on-prem) hosts. To give some sense of familiarity, only the three clouds with the lowest exposure rates bear labels. The point is not whether Cloud A is "better" than Cloud B, but rather that substantial variation exists among them. We cannot explain why the provider at the top of the list has an exposure rate 144× that of Oracle, but we suspect it has a lot to do with the nature of those clouds and how they're used. Perhaps systems in Oracle's cloud primarily host major enterprise applications that are rigorously maintained by their owners. Perhaps the unnamed cloud on top plays home to a higher share of SMBs and/or test workloads. We simply don't know. But we can safely conclude that scattering your hosts randomly across cloud providers is unlikely to achieve positive outcomes. If you do go, "where?" is the next—and equally important—decision.

None of this discussion deals with common-mode failure among cloud suppliers such as the Meltdown [3] and Spectre [4] issues announced in January 2018. Rather, it asks a fuzzy question: what is the causal relationship here? Is it size? Is it diversity? Is it complexity in some other sense? Can the causal mechanism be identified and sufficiently well understood to drive policy? What more data would help (or would more data help)?

As with other budding romances, "Should I stay or should I go? (Don't you know which ~~clothes~~ clouds even fit me?)"

References

[1] D. Linthicum, "Cloud Complexity Management (CCM): A New Year, a New Problem": https://www2.deloitte.com/us/en/pages/consulting/articles/cloud-complexity-management-a-new-year-a-new-problem.html.

[2] https://www.riskrecon.com/.

[3] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, M. Hamburg, "Meltdown: Reading Kernel Memory from User Space," in *Proceedings of the 27th USENIX Security Symposium (USENIX Security '18)*, pp. 973–990: https://www.usenix.org/conference/usenixsecurity18/presentation/lipp.

[4] P. Kocher, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, Y. Yarom, "Spectre Attacks: Exploiting Speculative Execution": https://arxiv.org/abs/1801.01203.