INTERNET RISK SURFACE REPORT

Exposure in a Hyper-Connected World

A collaborative research project between RiskRecon and the Cyentia Institute

riskrecon™

119 Cyentia INSTITUTE

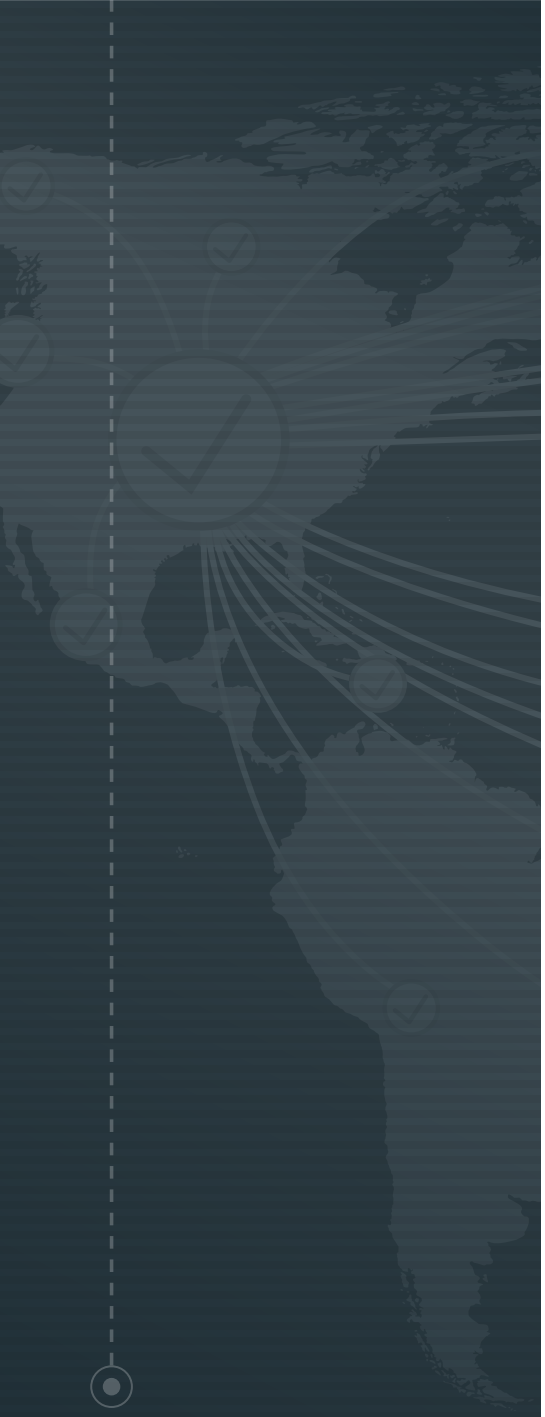

This research was commissioned by RiskRecon.

RiskRecon collected the dataset and provided it to the Cyentia Institute for independent analysis and drafting of this report.

Table of Contents

Some Opening Thoughts

FROM RICHARD SEIERSEN

RICHARD SEIERSEN IS THE CEO OF WWW.SOLUBLE.AI AND THE AUTHOR OF "HOW TO MEASURE ANYTHING IN CYBERSECURITY RISK" AND THE FORTHCOMING "THE METRICS MANIFESTO: CONFRONTING SECURITY WITH DATA"

Soluble

"Risk is when you make a bet, and if you get it wrong, it may hurt...maybe a lot.

Chevelier de Mere (1607-1684) was a "gentleman" gambler and proxy innovator. I say gentlemen with quotes because the title of Chevelier was self prescribed. And I say proxy innovator because he is quasi credited with creating the field of risk management and probability theory.[1] What was the source of his inspiration? He was tired of losing at dice. De Mere intuited that there must be some underlying property in dice he could exploit to his advantage. To get help he engaged his friend Blaise Pascal , arguably the top mathematician of his day. His intuition worked! De Mere started to win thanks to Pascal. And mankind won too! We got probability theory and perhaps an early nudge into risk management.

De Mere's story is a bit like Billy Beane's. Beane is the general manager of the Oakland A's and the protagonist of the great book and movie MoneyBall. Like de Mere, he also wanted to lose less. He intuited there must be some way to beat the odds so he also enlisted the help of an elite card counter—or shall we say statistician. They found that just getting on first base, any way and any how, correlated highly with winning. And as they focused their players on getting on first base winning indeed ensued!

Those of us attracted to risk management are perhaps a bit like Billy Beane and de Meres. We intuit there must be some property to the game we are playing that can give us an edge. I know I think this way. It's why I wrote my first book and am nearing completion on my second. But you may argue, "security is not dice and it's not baseball. We have intelligent adversaries and N parties playing with our data in parts unknown!"

Indeed, the rules are not as simple as dice and the surface we play on is far more multi-dimensional than a baseball diamond. Should we throw our hands up in despair because our game is complex? Or, should we try to understand the rules and surface of our game?

The research report you are about to read is on the latter topic, something RiskRecon calls "Risk Surface." It's made of all the digital things you make bets on, and if you get it wrong, it hurts! A good example is the SaaS companies you pump customer data into as well as all their cloud providers. RiskRecon is in the business of making that "Risk Surface" visible. Like de Mere and Bean, the question for you is, "can we lose less (and by converse win more) by understanding, measuring and managing this risk surface?" I say yes!

I hope you enjoy this critical piece of risk surface research ensembled by RiskRecon and the brilliant team of researchers at the Cyentia Institute. And I hope it leads you to winning more, or perhaps losing less, than your competitors. However you choose to frame the outcome - you are absolutely playing this game. You might as well try to move the odds in your favor.

[1] Cooper, Dan, and Brian Grinder. "Probability, Gambling and the Origins of Risk Management." Financial History Winter (2009).

65%

of hosts sit on an external network; 27% of firms host assets with at least 10 external providers.

84%

of organizations host critical or sensitive assets with 3rd parties.

35%

of firms have externally-hosted assets with high or critical findings.

Key Findings

The Digital Transformation era ushered in many operational and strategic benefits for modern organizations. It also fundamentally changed our dependence upon the internet and a myriad of interconnected 3rd and 4th parties for key business activities. And with these dependencies come new risks. Exploring and measuring the resulting "internet risk surface" is the purpose of this report. It is the first offering from an ongoing research initiative between RiskRecon and the Cyentia Institute.

In the pages that follow, we share our analysis of a fascinating dataset spanning millions of internet-facing hosts from thousands of firms and major hosting providers around the world. Here's a sampling of what we uncovered.

1. **An organization's internet surface area is larger and more complex than you might think:**
   - The typical organization has 22 internet-facing assets (or "hosts"), but some maintain over 100,000.
   - 65% of hosts sit on infrastructure owned by an external entity.
   - 27% of firms host assets with at least 10 external providers.
   - 57% of organizations have hosts in multiple countries; 6% spread across 10 or more countries.
2. **Organizations place a huge amount of trust and value in the hands of external service providers:**
   - 20% of an organization's internet-facing assets have highly sensitive data or functions.
   - 61% of these high-value assets are hosted on 3rd party networks.
   - 84% of firms host critical and/or sensitive assets with 3rd parties.
3. **Exposures exist in all areas of the internet risk surface, but appear to aggregate in 3rd party networks:**
   - 56% of organizations exhibit severe findings in at least one asset.
   - 35% of firms have high or critical findings in assets hosted with external service providers.
   - Overall, organizations are 3X as likely to have high-value assets with severe findings off-prem vs. on-prem.
   - But over half of firms actually show fewer security findings in assets hosted on external infrastructure.

CHAPTER 1

Dataset Firmographics

For the purpose of this research, RiskRecon provided Cyentia a large anonymized sample of their 3rd party risk assessment database. The sample is representative of the organizations in their production dataset. It contains sanitized information on 18,000 organizations and more than 5 million hosts located in 200+ countries. Across those hosts, RiskRecon identified over 32 million security findings of varying severity.

A breakdown of the industries represented in this report is found in Figure 1. We base these on the top-level sectors (or groupings of sectors) as defined in the North American Industry Classification System (NAICS). If you're interested in determining a NAICS sector for a particular entity or type of entity, this tool will help.

FIGURE 1: Industry representation by percentage of organizations and hosts

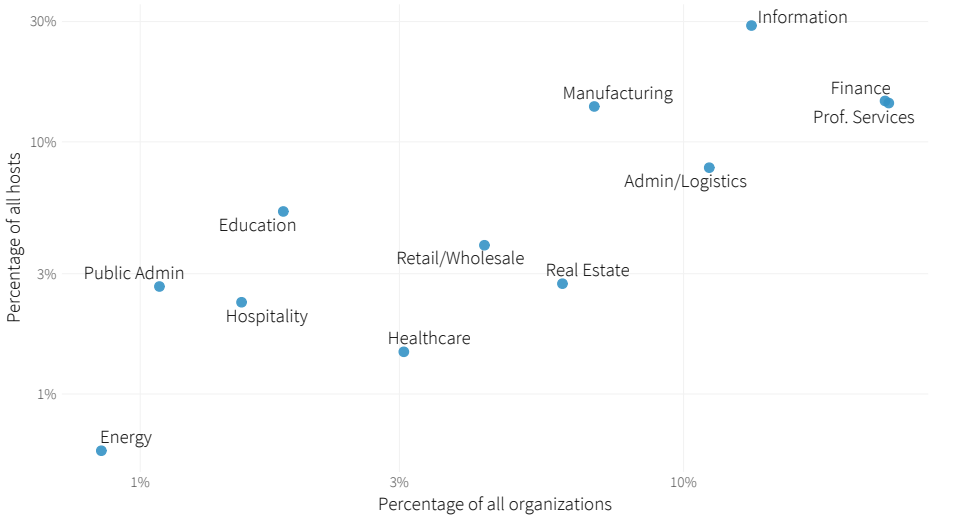

The orientation of the sectors from left to right in Figure 1 corresponds to the relative frequency of organizations in that sector. Alignment from top to bottom compares the proportion of hosts (Internet-facing assets) across industries.

The orientation of the sectors from left to right corresponds to the relative frequency of organizations in that category. So, we have more firms from the Professional Services and Finance sectors than from Energy. Alignment from top to bottom compares the proportion of hosts across industries. Given its subsectors, it is not surprising that the Information sector boasts a clear lead on the host axis. We show both views here because charts throughout this report use one or the other, and this should help calibrate interpretations.

FIGURE 2: Organization size by employee count

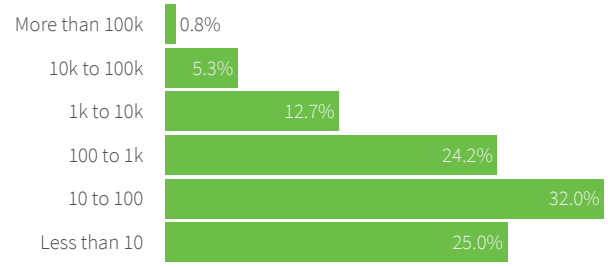

FIGURE 3: Organization size by annual revenues

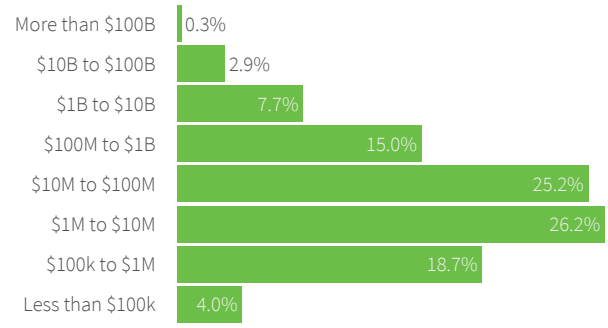

FIGURE 4: Organization location by region of operation

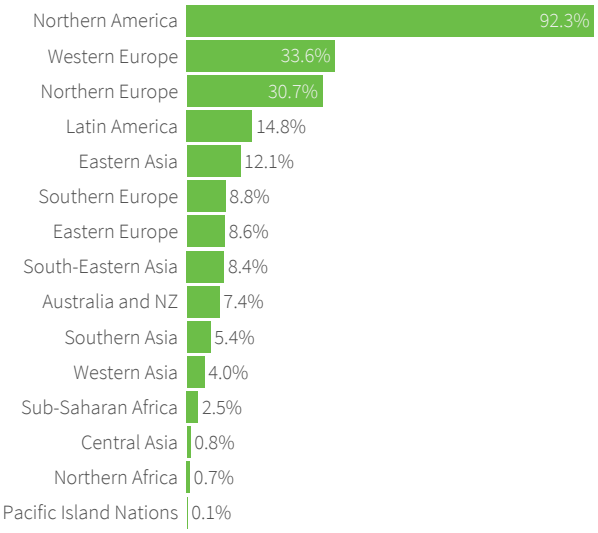

We have several ways of viewing organization size from the dataset. Figures 2 and 3 cover two of those—number of employees and annual revenue. It's apparent from either measure that the sample skews toward smaller and midsize firms. But the same is true among all registered companies; therefore, this skewing reflects the broader population. If anything, larger organizations are statistically over-represented in this sample.

Regionally, organizations with a presence in Northern America and Europe dominate the dataset. But Latin America, Australia and New Zealand, and Eastern/Southern Asian subregions have decent proportional representation as well. Coverage across Africa and in the Pacific Island Nations is fairly sparse, so organizations in those regions should take that into consideration.

DATA COLLECTION

Riskrecon gathers and collates data from a wide array of sources to build a comprehensive view of organizational risk. The process starts by identifying a target organization for analysis. Firmographic data is collected from 3rd party sources to establish the organization's size, revenue, history, and geographic locale. Along with this information, seed intelligence in the form of domains and netblock ownership is established. This information is then used to expand the list of domains and hosts associated with an firm. Hosts are scanned to establish any services they make available as well as what type of software they may be running. From this information, Riskrecon is able to establish what vulnerabilities or safety measures are present and make inferences about the value of the asset.

CHAPTER 2

Measuring Internet Risk Surface

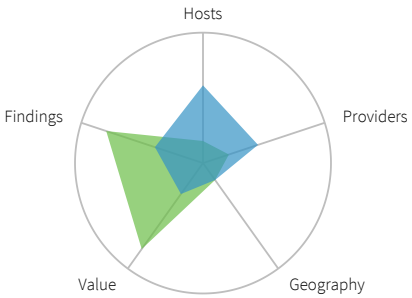

Risk surface refers to anywhere an organization's ability to operate, reputation, assets, legal obligations, or regulatory compliance is at risk. The aspects of a firm's risk exposure that are associated with or observable from the internet can be considered its internet risk surface. Since a huge portion of a modern organization's value-generating activities relies on internet-enabled processes and 3rd party relationships, that surface is much more extensive than one might expect. In this section, we identify and measure key aspects of the internet risk surface through the data sample collected by RiskRecon.

> This section builds toward a view of risk surface shown in the figure above. But first we need to understand the individual dimension of that surface.

Key Measures of the Internet Risk Surface

One could identify numerous inputs of potential use in measuring an organization's internet risk surface based on the broad definition above. After reviewing the available data, we selected five key measures to give structure to our exploratory analysis:

- **Hosts**: Number of internet-facing assets associated with an organization
- **Providers**: Number of external service providers used across hosts
- **Geography**: Measure of the geographic distribution of a firm's hosts
- **Value**: Relative sensitivity and criticality of hosts based on multiple indicators
- **Findings**: Security-relevant issues that expose hosts to various threats

We give individual attention to these measures in the subsections that follow and then bring them together in an example contrasting two different companies. After that, we provide large-scale comparisons of internet risk surface across firmographic segments.

Number of Hosts

This is not exactly breaking news, but some organizations have more internet-facing hosts than others. Figure 5 provides some data to expand on that truism. The median number of hosts per organization is 22, but the long tail shows some with more than 100,000 under their purview. That matters because protecting a large internet presence is a different ballgame than protecting a tiny one, regardless of any other factors.

We often think of organizations as "large" or "small" based on employee count or annual revenue, but how do such size descriptors apply to a firm's internet footprint? Figure 6 confirms that a correlation among them does indeed exist. In general, we see purple dots (representing firms with less than 10 hosts) toward the lower left and yellow and green dots toward the upper right (firms with 1000+ hosts). But we can also discern that not every organization fits that mold due to the substantial intermingling of colors along the spectrum.[2] Thus, we find organization size and internet surface to be related but still worthy of measuring separately.

FIGURE 5: Distribution for the number of hosts per organization

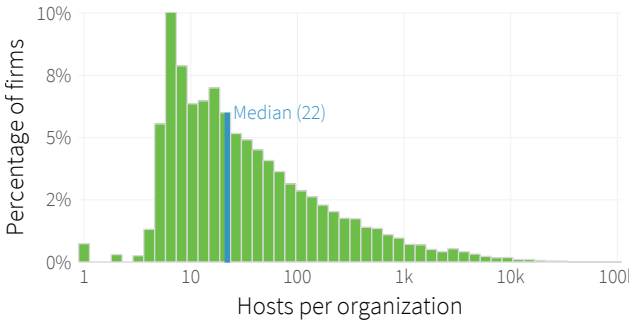

FIGURE 6: Relationship between employee count, revenue, and internet hosts. In general, we see purple dots (representing firms with less than 10 hosts) toward the lower left and yellow and green dots toward the upper right (firms with 1000+ hosts).

[2] We couldn't resist giving props to the lonely yellow dot firm toward the upper left banking $10B in revenue with less than 100 employees but a massive number of hosts. Sounds like a botnet biz model.

EXTERNAL PROVIDERS

In addition to the number of hosts under management, their placement inside or outside organizational infrastructure is very important in shaping the risk surface. There was a time when firms could walk over and point to their IT assets, but that time has long passed. The IT footprint of modern organizations tends to be undefined and highly distributed across a plethora of external service providers that own, control, or manage assets. That matters because a lack of clear visibility into all assets—wherever they're hosted—means a lack of visibility into a firm's true risk posture.

FIGURE 7: Proportion of hosts on external infrastructure per firm

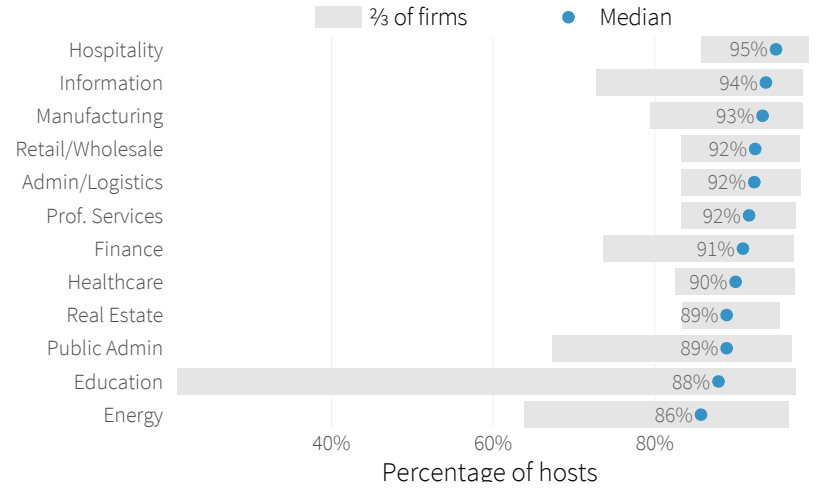

The notion of internal vs. external hosts can be viewed from several angles. At the most fundamental level, RiskRecon determines whether the asset resides on a netblock owned by the organization (internal) or another organization (external). Figure 7 establishes a rather high proportion of external assets for all industries, and Figure 8 shows how this varies by organization size. It basically confirms what we already know: small firms own little of their infrastructure, but that changes with growth (though it rarely tips to an internal majority).

FIGURE 8: Proportion of hosts on external infrastructure by organization size

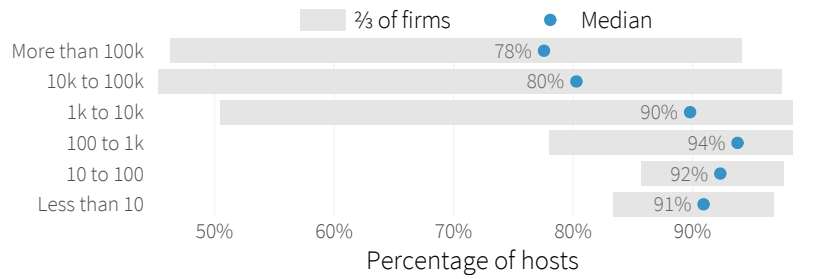

A lack of clear visibility into all assets–wherever they're hosted–means a lack of visibility into a firm's true risk posture.

Figure 7 is the first of many charts in this format. The blue dot marks either the mean or median value. The grey bars encompass the middle 2/3rds of firms in each segment. We do this to give a sense of what's "typical" (the mean or median) for the organizations represented as well as the amount of variation (gray bars) among them.

Figure 7 is a great example of where this is helpful. The median percentage of externally-hosted assets is high for all sectors, but the gray bar for Education is unusually wide. This indicates many institutions vary substantially from the norm (i.e., host more assets internally).

65%

of hosts reside on an external network.

FIGURE 9: Top external providers by proportion of hosts

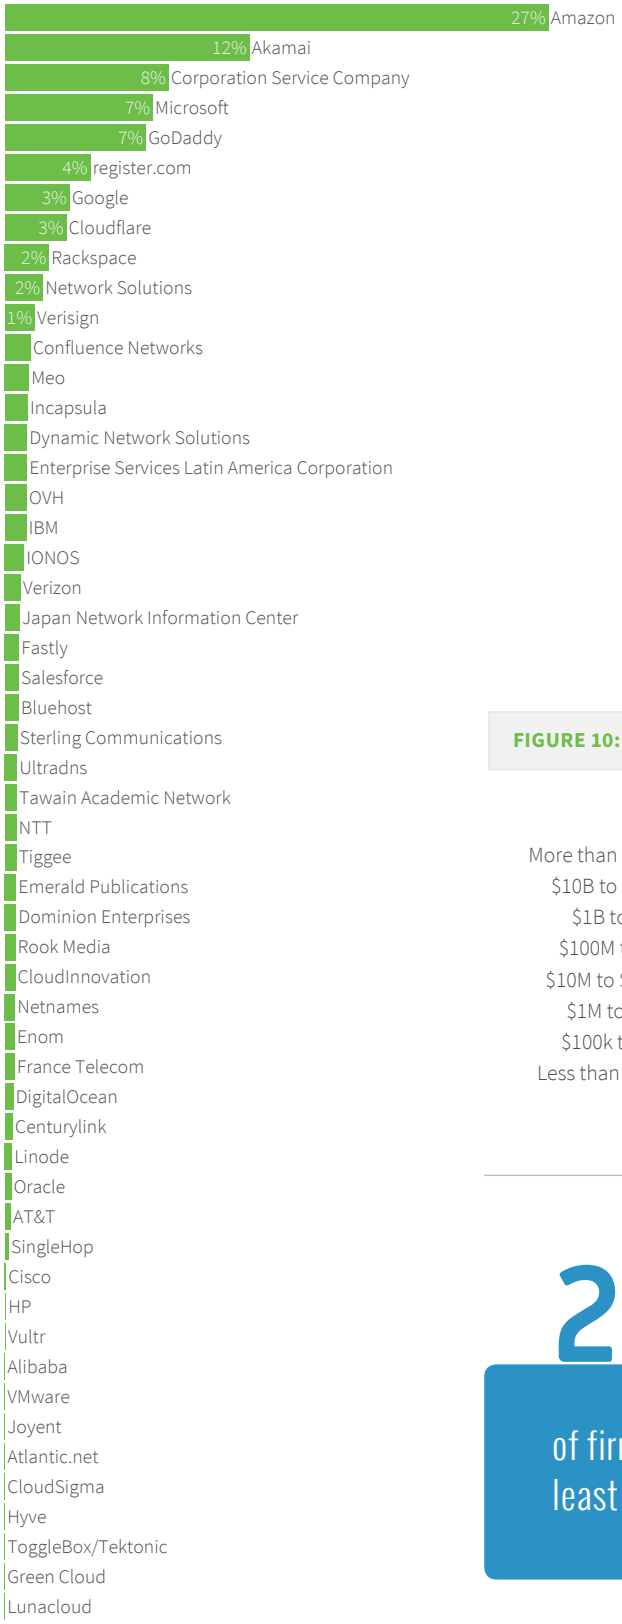

For any external host, RiskRecon determines and records who owns the infrastructure. Figure 9 lists the top 50 external providers based on the number of associated hosts. The list represents a mix of cloud providers, content delivery networks, DNS, telecommunications services, etc. We keep things simple by sticking to the overall internal vs. external distinction in this discussion of risk surface measures, but it's possible to analyze a subcategory of particular interest. In fact, we will do exactly this in a future report focusing on cloud providers.

Looking over Figure 9 likely reveals more than one of your own providers. Turns out you're not alone in that regard, especially if hailing from a larger firm. Figure 10 demonstrates that the number of external service providers grows dramatically with (and probably to support) revenue. This corroborates our earlier claim that the risk surface is highly dependent upon multiple 3rd parties.

FIGURE 10: Number of external providers by firm revenue

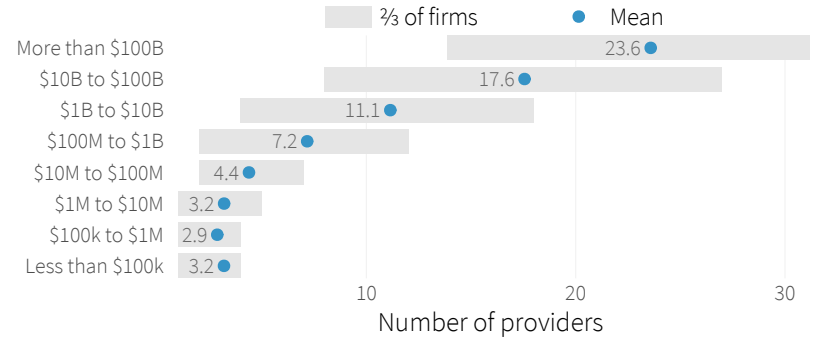

27%

of firms host assets with at least 10 external providers.

84%

of firms host high-value assets externally.

GEOGRAPHIC DISTRIBUTION

We often talk about the internet as borderless, with little regard for physical geographies. While the virtual and physical worlds differ in many ways, it's a simple fact that every device on the internet resides (or moves) somewhere in the real world. We consider this an important aspect of the risk surface because those different geographies have different policies, regulations, and customs that govern hosts and data. Thus, organizations with larger geographic footprints must manage a larger portfolio of geopolitical, legal, compliance, and physical risks tied to those geographies.

As an indicator of those complexities, we can examine the proportion of hosts located within and outside of a firm's home country of operation. Figure 11 provides this distinction. The two spires on opposing ends tell us that many organizations host all of their assets on foreign soil (left spike), many keep it all domestic (right spike), and many maintain a ratio between those extremes. Overall, the distribution leans toward larger proportions of domestic assets.

FIGURE 11: Hosts located in organization's home country

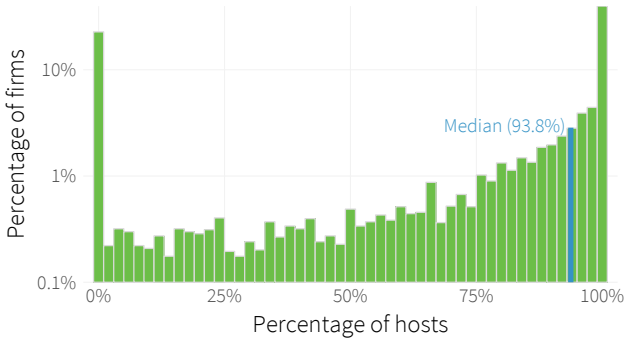

You might infer that this geographic dispersion of hosts differs by firmographic factors, and indeed it does. Not surprisingly, larger organizations (employee count and revenue) tend to have a presence in more countries. Industries show variation too, but we find region offers the most interesting view. Yes, 'regional distribution by region' seems awkward, but what we're really viewing in Figure 12 is a mix of tendencies, tolerances, and necessities among regions regarding the geographic dispersion of information assets.

FIGURE 12: Geographic diversity of hosts by regional firms

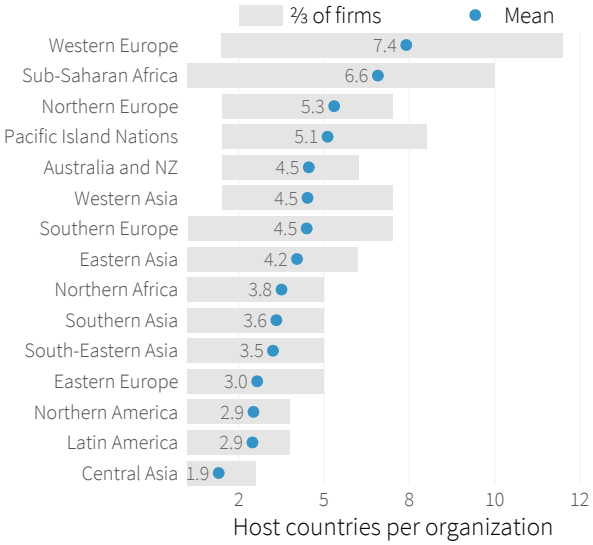

Some findings in Figure 12 are intuitive; others less so. The position of the United States near the bottom makes sense; so many of the major hosting and IT service providers (see Figure 9) hail from that region. Western Europe at the top falls in that "less so" bracket. Due to the stricter data sovereignty and privacy laws common to countries in that region, we tend to think of them as digital islands. But a firm's internet footprint covers a lot more than data storage. The European Union offers favorable trade conditions among member nations, which likely contributes to results in Figure 12.

57%

of a firms have hosts in multiple countries. 6% span 10 countries.

ASSET VALUE

Asset value is an integral component in RiskRecon's assessment and prioritization algorithms, enabling firms to create risk action plans that go beyond a simple list of issues to fix. Each host in the dataset has a value of low, medium, or high as determined by RiskRecon. In general, high-value assets collect sensitive information, authenticate users, run critical services, etc. Hosts assigned a medium valuation don't appear to perform such functions, but they're network neighbors to those that do. This makes them ideal pivot points into sensitive and critical environments. Static assets that aren't connected to higher-value systems fall in the low range.

Ternary plots such as Figure 13 work well for depicting ratios from three variables that comprise a whole. You can read more on interpreting them precisely here, but we'll get you to "good enough" in the next paragraph.

The position of each hex in the triangle represents a different proportional mix of assets categorized as low, medium, and high value. The shading corresponds to the number of firms represented at that position. Firms in the middle have equal balance among asset values. Those in the lower left corner have mostly low-value assets. Going up adds more and more in the medium category, and sliding to the right raises the proportion of high-value assets.

Keeping interpretation simple, the density toward the lower left indicates that asset value ratios cluster in the low and medium spectrum in the majority of organizations. But it is readily apparent that firms exhibit a wide range of different valuation mixtures from relatively balanced to heavy in one category.

FIGURE 13: Relative proportion of low, medium, and high-value assets per organization

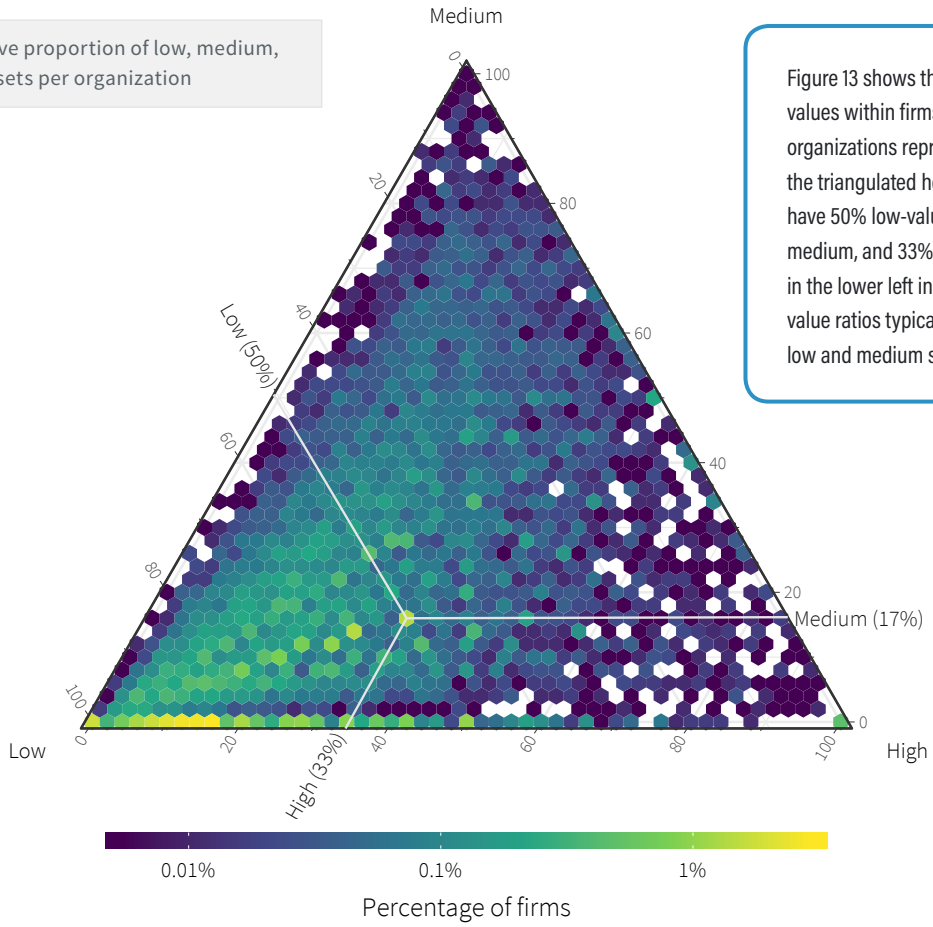

Figure 13 shows the "mix" of asset values within firms. For example, organizations represented by the triangulated hex in the plot have 50% low-value assets, 17% medium, and 33% high. The density in the lower left indicates asset value ratios typically cluster in the low and medium spectrum.

20%

of a firm's internet-facing hosts have highly sensitive data or functions.

The proportion of high-value assets in the Public sector is nearly

5x

that of the Hospitality sector.

Figure 14 emphasizes this relative balance among value tiers, with an added twist to compare that ratio across industries. The sectors on the left are sorted from top to bottom in terms of their proportion of high-value assets. Given that the Public sector exists to serve a diverse constituency, it's not terribly surprising to learn they maintain many sensitive and critical assets. Real Estate in the #2 spot seems rather odd, but think beyond your local realtor. That NAICS sector covers a wide variety of appraisal, brokerage, property management, and rental services that all require collection of sensitive information.

FIGURE 14: Asset value tiers by industry. Sorted by percent of high-value assets.

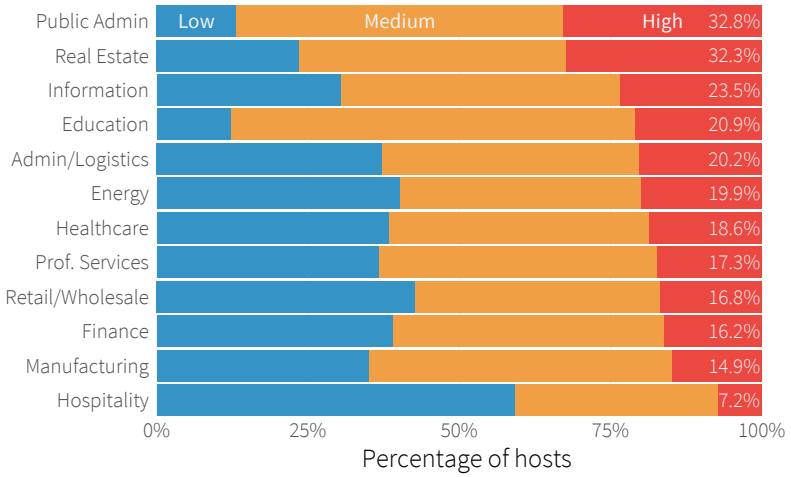

The position of Finance toward the bottom may also strike some as odd. Keep in mind that this does not suggest financial firms have inherently low-value assets; it simply shows the relative proportion of value categories among hosts within each industry. Many financial firms work hard to consolidate critical functions and assets to better manage them. At the same time, they have a huge amount of marketing material on the web. The Hospitality sector has a similar marketing-heavy web presence, which undoubtedly drives their industry-leading proportion of low-value assets.

We also examined asset value by employees and revenue. In general, the proportion of high-value assets decreases as organization size increases. We felt this worthy of mention because while small firms may not have many assets, those they have are critical to their bottom line.

> While small firms may not have many assets, those they have are critical to their bottom line.

Before leaving the topic of asset value, we thought it worthwhile to pause and reflect on what we've learned so far about the internet risk surface. We know that it varies significantly among organizations in scale across hosts, service providers, and geographies. But we haven't yet considered how those dimensions interrelate to one another. Figure 15 serves up some food for thought to that end.

Starting on the top of Figure 15, we begin with all hosts in our sample dataset (millions). The colored streams separate proportionally from there into the 3 asset value categories. From there, hosts distribute to internal (owned by the firm) and external infrastructure. Medium-value hosts split fairly evenly, but less valuable assets are much more likely to be hosted with a service provider. As just one example of this, consider the millions of brochure websites hosted with Wordpress. Why not let someone else deal with that—especially when they do it cheaper with less hassle?

Given the pattern of lesser tiers, one might expect critical and/or sensitive assets to remain largely on internal infrastructure. Figure 15 refutes that hunch and administers a dose of reality about the modern internet risk surface: 61% of high-value assets are hosted externally.[3]

The lowest tributaries in Figure 15 follow a geographic dispersion. Hosts located in the firm's home country are considered domestic, else they receive the foreign classification. We'll restrict commentary to the high-value assets, but feel free to follow the flows that interest you most. The terminal points of the red streams indicate organizations host nearly two-thirds of valuable assets domestically.[4] But the remaining one-third of high-value assets that reside in external infrastructure and/or foreign soil reminds us that we cannot assume the crown jewels are always kept near and dear.

FIGURE 15: Breakdown of asset value across internal vs. external and foreign vs. domestic infrastructure. Starts with all hosts in the dataset and provides proportional comparison at each level.

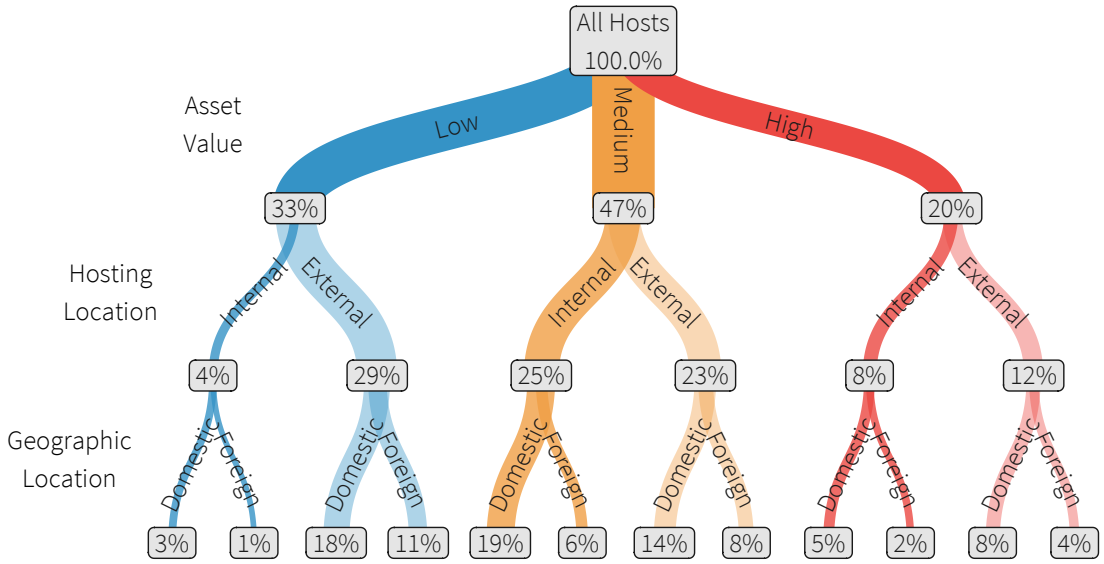

[3] 12% / 20% = 61% of high-value assets hosted externally

[4] (5% + 8%) / 20% = 65% of high-value assets hosted domestically

SECURITY FINDINGS

Last, but certainly not the least on our list of risk surface measures is an assessment of the security status and configuration of hosts. RiskRecon identifies a wide range of potential security issues, records the details of those findings, and assigns a severity classification based on several factors. While any finding could contribute to a costly security incident, we focus our analysis here on those deemed to be of high and critical severity.

FIGURE 16: Hosts with high or critical security findings per firm

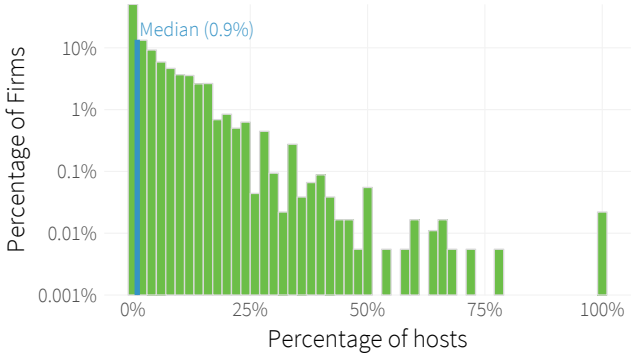

Overall, we found that about 1% of a firm's internet-facing hosts have at least one high or critical exposure. The majority range between 0% and 6%, but a small minority exhibit greater than 50% vulnerable hosts! We could measure exactly how many findings per firm or analyze the details of those findings (and will do so in future work), but the simple truth is that threat actors need just one vulnerable host to compromise an environment. And if we good guys can find that host, the bad guys can too.

Let's continue with an industry-based view. At the two extremes, Figure 17 confirms what many expect: the Public, Education, and Healthcare sectors have the highest average prevalence of severe findings and Finance with the lowest. But let's be careful not to let confirmation bias lead us too far astray. Keep in mind that differences across the board range only by a few percentage points and variation among firms (even in the same sector) is much larger than among sectors.

Figure 18 captures the prevalence of high and critical findings among organizations of different sizes. The results slam SMBs with a double whammy. Not only do they more than double the rate of security findings seen in larger firms, but those rates vary widely from one to another.

FIGURE 17: Hosts with high or critical findings by industry

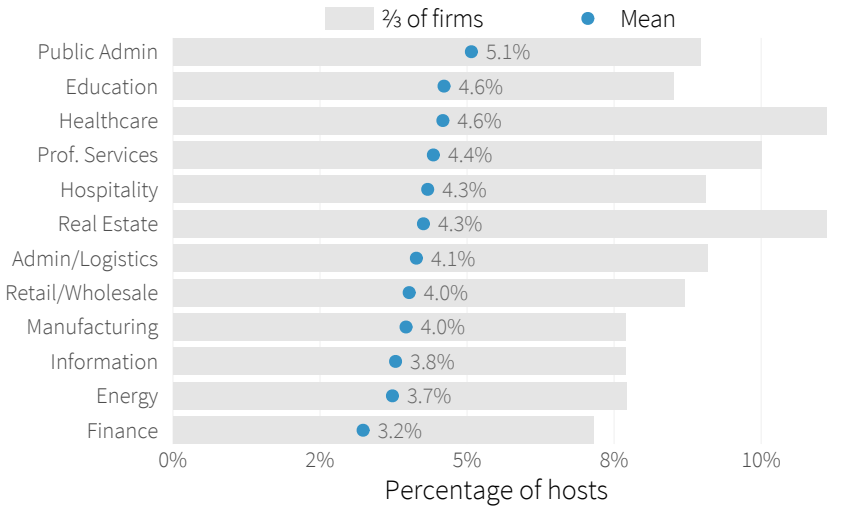

56%

of firms exhibit severe findings in at least one asset.

35%

of firms have severe findings in assets hosted with external service providers

Variation between regions ranges less than 4% from top to bottom, but those few points represent a nearly 4X jump in exposures when comparing Northern America and Western Europe to Eastern Asia!

Further note how the gray bars shrink steadily as employee count grows, indicating less variation in the proportion of severe findings among larger firms. In other words, the bottom half of SMBs are much worse off than underperforming enterprises. That effectively translates to more risk. So, choose wisely if SMBs have critical roles in your digital ecosystem.

FIGURE 18: Hosts with high or critical findings by organization size

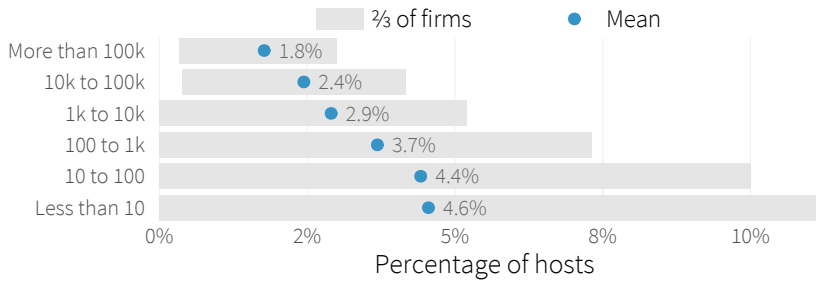

Moving on, Figure 19 compares the proportion of hosts with high and critical findings across regions. Hosts located in Eastern Asia exhibit the highest prevalence of findings, but even the regions with the lowest rates can't boast too much. Overall, variation between regions ranges less than 4% from top to bottom. Those few percentage points, however, represent a nearly 4X jump in exposures when comparing Northern America and Western Europe to Eastern Asia! And that might be enough to influence the decision criteria when sourcing partners or deploying services.

FIGURE 19: Percentage of hosts with high or critical findings

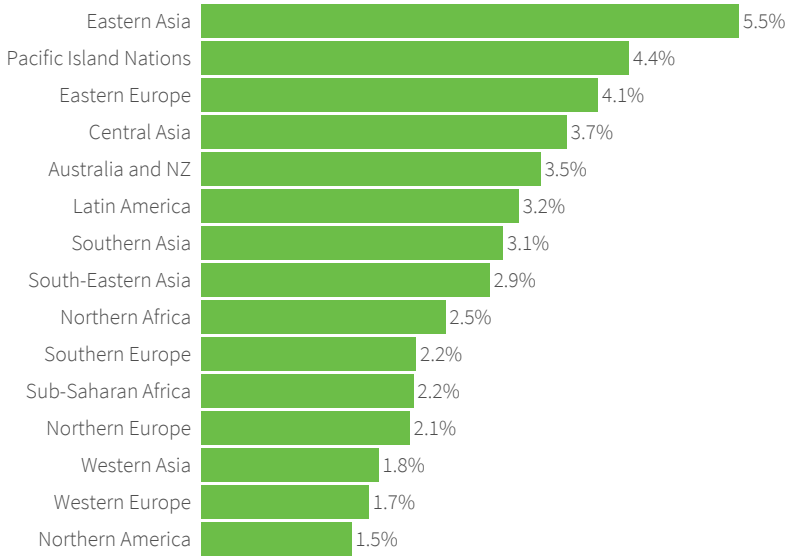

Overall, firms are

3x

as likely to have severe findings off-prem vs. on-prem.

"

We could interpret these results to claim external service providers must be incompetent or negligent. But Figure 20 may speak more to the capabilities of the asset owner than the hosting provider.

Beyond geography, we want to know how the prevalence of severe issues among hosts on premises compares to that of hosts in external service providers. Figure 20 supplies that comparison. Each square in Figure 20 represents 1% of organizations in the sample dataset meeting the criteria for this analysis. Those firms host at least one asset in their own infrastructure (top row) and that of external providers' (bottom row). At the leftmost tier, there's little difference because most host something on and off premises.

Things get interesting with the middle and right tiers. We see that 88% of firms host high-value assets on their own networks, while 57% trust some of their most valuable systems with 3rd parties.[5] Proceeding to the right reveals the shape those assets are in security-wise. A somewhat modest 8% of firms exhibit high or critical findings in a high-value asset hosted on premises. For organizations using external providers, that statistic jumps threefold to 25%!

We could interpret these results to claim external service providers must be incompetent or negligent. But keep in mind those external providers often aren't responsible for securing hosts you place in their infrastructure. So Figure 20 may speak more to the capabilities of the asset owner than the hosting provider. But regardless of who's to blame (we suspect both parties share it), the fact remains that organizations are more likely to have a high-value assets with severe issues hosted externally than they are for those hosted internally. And that's why we assert the need for measuring the complex and interwoven dimensions of the internet risk surface.

FIGURE 20: Firms with high or critical findings in internal vs. external infrastructure

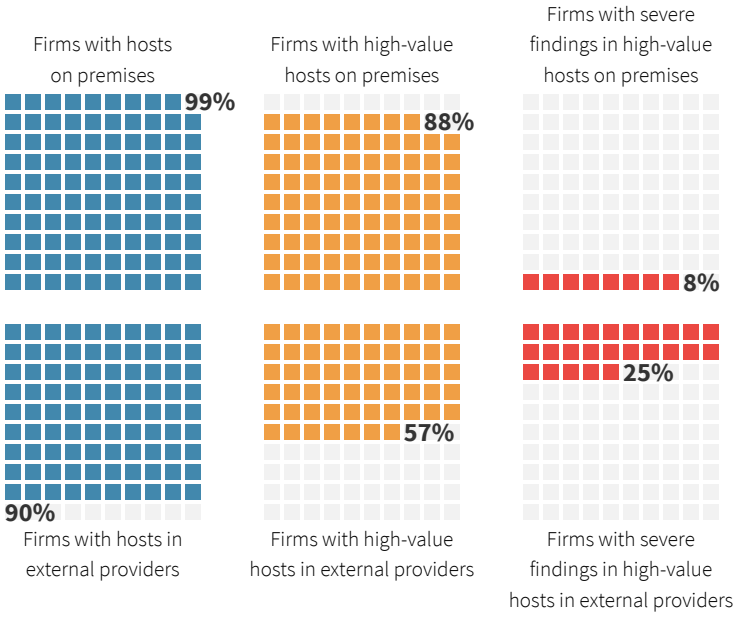

[5] This statistic differs from Figure X because we're looking at organizations here rather than hosts.

CHAPTER 3

Comparing Internet Risk Surfaces

We've examined these five measures separately to this point, but in reality, they collectively describe an organization's internet risk surface. We now present them holistically in an example that compares two retail firms that share similar demographic profiles. After that, we provide large-scale comparisons of internet risk surface across firmographic segments.

A Tale of Two Retailers

> Having a single view of all these factors is critical to any proper assessment of internet risk surface.

Both retailers selected for this example are based in the same country with annual revenues between $400M and $500M and 1000 to 2000 employees. Figure 21 compares the key measures of their internet risk surfaces and shows their very different risk profiles.

Figure 21 takes the form of a standard radar (aka "spider") chart, with a spoke for each of the five risk surface measures examined in the previous sections. The position of points along those spokes marks the firm's relative value for that measure among all organizations. The closer a point is to the perimeter, the closer that firm is to the maximum value among other firms for that measure. The two polygons formed around those points represent the internet risk surface of the two retailers. In general, the more area the polygon covers, the larger the risk surface the firm needs to manage.

With that explanation in hand, it's easy to see that what sets the green retailer apart is higher than normal asset values and fairly high-levels of security findings. All of

that value and exposure concentrates in relatively small number of hosts and hosting providers. The blue retailer exhibits a comparatively bigger internet footprint and relies on a larger number of external service providers. Despite this, it manages to keep security findings comparatively low. Both fall well below the norm in geographic diversity among hosts.

FIGURE 21: Internet risk surface charts for two similar retailers

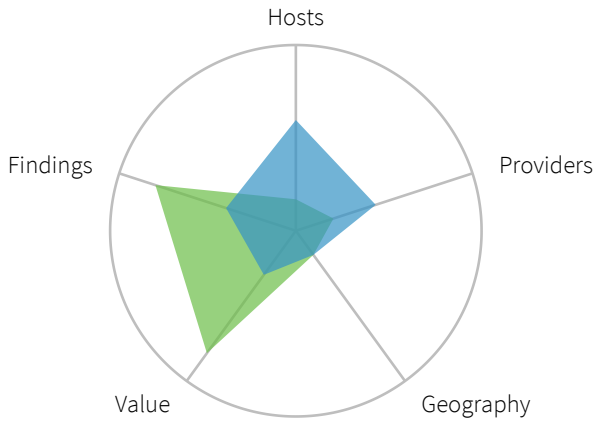

We've already established that firmographics and risk surface are not always in sync, so contrasting two similar firms in this way may not blow any minds. But it should reinforce the point that having a single view of all these factors is critical to any proper assessment of internet risk surface. From a practical perspective, Figure 21 poses an interesting question for 3rd party risk management—which retailer would you do business with?

Assume for a moment you represent a growing manufacturer looking for a trusted partner to bring your goods to market. Further assume you'll integrate IT systems with the chosen retailer to coordinate logistics and share inventory and sales data. Blue Retailer has been doing this for a long time and would provision a solution via multiple 3rd parties. Green Retailer is younger, more technologically innovative, and leverages that trait to offer you a direct solution at a better price. Who do you choose and why?

Obviously, there's no definitively right or wrong answer here. Ultimately, it's a business and risk-based decision. Maybe Blue Retailer's reliance on 3rd parties concerns you. Maybe you don't want the exposure presented by Green Retailer's higher rate of findings. What should be obvious is that having information like what you see in Figure 21 would help ensure you had all relevant facts on the table during the decision-making process.

> "It's easy to see that what sets the green retailer apart is higher than normal asset values and fairly widespread security findings."

> "It should be obvious that having information like what you see provided in Figure 21 would help ensure you had all relevant facts on the table during the decision-making process."

Risk Surface Profiles

This subsection incorporates the individual measures of internet risk surface into a collective view for comparison across size, industry, and region. If you skipped the *Tale of Two Retailers* piece, we strongly suggest reading that first. The format of these charts is the same, and so we won't rehash how to interpret them. The only difference here is that each radar chart now represents many firms within the designated segment. The extent of each segment is the median value among organizations.

Given so many dimensions, we could make almost limitless observations and speculations concerning these charts. But we suspect each reader will have something different in mind, and so we'll simply highlight a few general patterns to watch for as well as a few specifics that caught our eye. Let's start with the risk surface charts comparing organizations across revenue tiers.

Results across revenue segments in Figure 23 demonstrate a progressive pattern and seem to tell a familiar story. Small firms appear as a sliver, with only a few high-value assets and not much else. Revenues grow and so do the number of hosts and security issues affecting them. As their internet footprint continues to expand, they rely on more and more external service providers. Eventually, they regain control and the prevalence of findings begins to shrink.

FIGURE 22: Internet risk surface charts comparing firm size in revenue

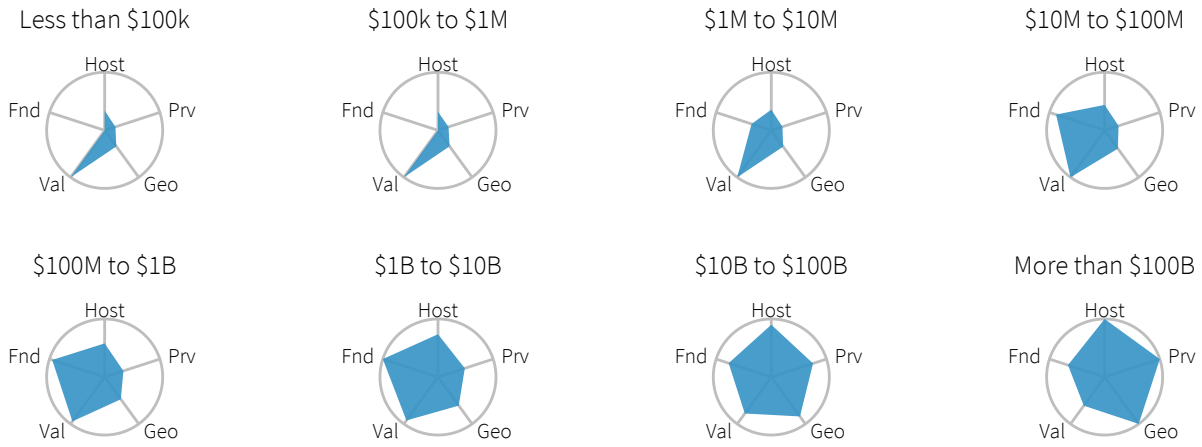

Abbreviations in Figures 22-24:
- Host = Hosts
- Prv = Providers
- Geo = Geography
- Val = Value
- Fnd = Findings

> As revenues grow, so do the number of hosts and security issues affecting them.

Looking at Figure 22, we're drawn to the different shapes and sizes exhibited among industries. Education and to a lesser extent, Manufacturing, certainly stand out with large risk surface areas. Universities often have balkanized IT infrastructures leading to a sprawling footprint that's difficult to control.

The Finance sector takes on a unique shape. A low rate of severe security findings relative to higher positions along the other axes forms a noticeable "dent" in their risk surface. We find that metaphor rather appropriate; they seem to be making a dent in managing an otherwise large risk surface.

Similar risk surface patterns ostensibly indicate industries that may be connected in some way or share characteristics relative to the five key measures. For instance, the Hospitality and Retail sectors appear nearly identical. Could that be because their use of the internet is similar (e.g., heavily oriented toward customer service and payment processing) and IT is not typically their core competency?

FIGURE 23: Internet risk surface charts comparing industry sectors.

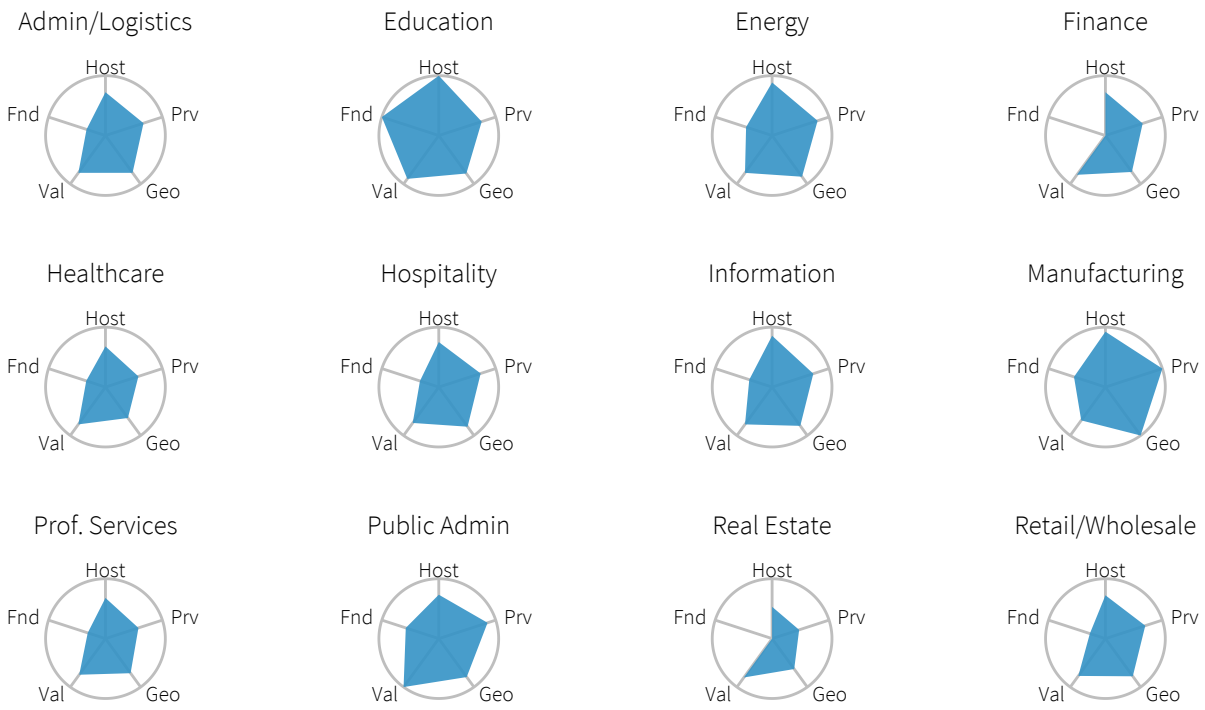

> Similar risk surface patterns ostensibly indicate organizations share characteristics.

This sounds like a dad joke, but the risk surface charts for regions in Figure 24 are all over the place. We struggle to find a coherent pattern or story, but perhaps that observation itself is noteworthy. Regions reflect the organizations of all types and sizes within them.

And that brings us to a final important point. Take care not to fall prey to the ecological fallacy and assume all organizations within a specific segment (say Australia) must share the same exact risk surface.

These figures represent the general trend among firms in each industry, revenue bracket, or region. As we learned in the *Tale of Two Retailers* section, similar firms can have very dissimilar profiles. And that's exactly why having visibility of your own internet risk surface—and that of your value chain partners—is so crucial.

FIGURE 24: Internet risk surface charts comparing regions

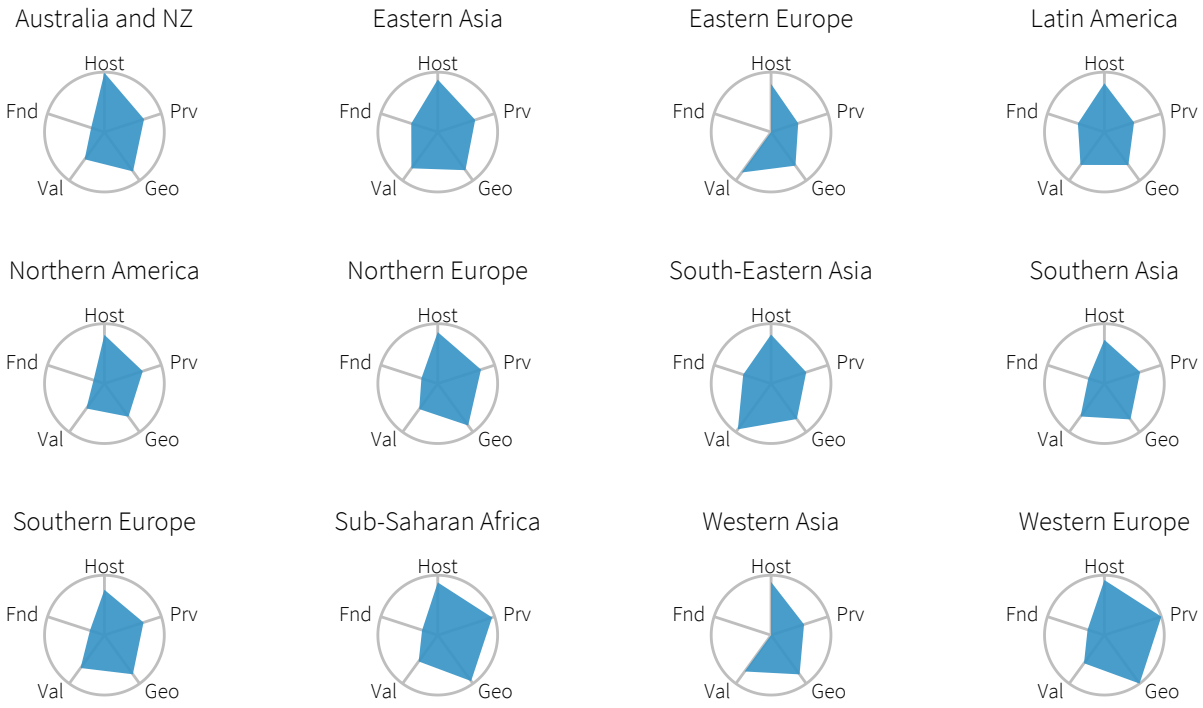

> Risk surface charts for regions reflect the organizations of all types and sizes within them.

CHAPTER 4

Conclusion & Future Work

Rather than the typical "tell 'em what you told 'em" format often adopted by conclusion sections, we'd like to end this report by beginning several more. RiskRecon and the Cyentia Institute will be working together over the next couple of years, and we're excited about the many research opportunities that lie ahead for this partnership.

One of the signs of good research is that it tends to prompt new questions through the process of answering others. We came into this project with the goal of exploring the internet risk surface—and we indeed learned a lot—but we leave it with a ton of nagging "what if we also looked at…" ideas running through our heads. We're content to address those in due time, but before signing off on this report, we wanted to at least peek at one more topic that begged for some attention.

> "There's a great deal of value at risk in IT service providers that are not ultimately responsible for owning the risk. But is that a bad thing?

We threw a lot of information and statistics at you detailing the widely-distributed and highly-dependent nature of the internet risk surface. After reading this report, it should be obvious that there's a great deal of value at risk spread across the infrastructure of IT service providers that are not ultimately responsible for owning the risk. But is that a bad thing? Often those providers make tools available to help secure hosts, which could be a good thing. Do we have any evidence that organizations are more or less capable of securing assets hosted on their own vs. somebody else's infrastructure?

Figure 23 gives us an inkling of an answer to that question. The horizontal and vertical axes represent the percentage of those hosts with severe issues internal or external infrastructure, respectively. Each point represents an organization, and color indicates whether that firm's hosts exhibit fewer severe exposures ("safer") on-prem or off-prem.

FIGURE 25: Comparison of hosts with severe findings in on-prem vs. off-prem infrastructure

At first glance, the question of which location is safer appears to be a toss-up. A little more staring and you may perceive a slight advantage on the blue side, suggesting more firms are better off with external providers. But your eyes can be deceiving; don't trust them! Important determinations like this should be adjudicated with statistical models designed to tease apart or reject subtle differences. The best conclusion we can draw from this figure is that some organizations do better with external hosting and some do better internally, and it's a pretty even split.

But of course this begs more questions. What types of firms do better on-prem vs. off-prem? What trends in performance can we discover across various firmographics? How do the services and role of a host affect the likelihood of severe issues? Do more issues translate to more breaches or higher losses?

We won't dive down the rabbit hole of statistical models, significant differences, fixed and random effects to answer these questions quite yet. That is another report (and another, and another...) for another day. What we've done here is explore and map the landscape of the internet risk surface. With this map in hand, we're more equipped to start digging in the right places to uncover a better understanding of cyber risk.

Thank you for taking the time to read this report.
Until next time...

> The best we can draw from this figure is that some organizations do better with external hosting and some do better internally, and it's a pretty even split.

A STUDY SPONSORED BY RISKRECON, INC