

November 3<sup>rd</sup>, 2023

Dear Partners,

As stocks fell, our deployment picked up starting in August and continuing in September and October. Keeping our hurdle rate high kept us from chasing at valuation highs in the summer and is now allowing us to engage with companies new and old at prices often 40-50% below their summer peaks. When a business with a great product in our sweet spot hits a valuation in our sweet spot, we buy, sometimes aggressively, depending on the magnitude of the undervaluation and the potential up/down asymmetry.

When we buy, it tends not to be at what turns out to be the absolute low; especially in this environment, our purchases have generally continued to decline after we've bought them. When that happens, as a good return turns into a very good and then great return, we've bought more, in incremental steps corresponding to the attractiveness of the risk/reward. Generally we've maxed out our position sizes at 8-9% of capital, but we're willing to go above that for investments we consider exceptionally certain and exceptionally attractive.

Could we have been more patient in buying? It's possible, but generally that would have required a clearer crystal ball for macro than the one we have with only the murkiest of glimpses. Not buying a very good asset that we know well and like and above our already-high IRR hurdle, in the hopes that the IRR goes up further, feels speculative. (Whereas we consider waiting for IRRs to go from bad to good to be not speculative, but just discipline. I think that is a distinction that investors who criticize holding cash sometimes miss, when some call out any cash holding as being market timing. Discipline in not investing our capital at bad IRRs has helped us repeatedly in the short period since launch.)

We are keenly aware of the risks involved in averaging down a position. Over the years we've read and re-read John Hempton's classic piece on this repeatedly, keen to avoid, as he puts it, "the value investor trap of losing 18 percent on a 7 percent position." When we average down, we focus on companies 1) without much financial leverage, and ideally with substantial net cash; 2) without significant operational leverage, in terms of margin risk; 3) without risk of technological obsolescence, which ties into our focus on great products; 4) with less-than-average exposure to discretionary purchases or commodity prices that would be heavily impacted by a recession; and 5) where we can anchor to some metric that provides solid valuation support not too far below the current price – e.g. in software, a laser focus on hypothetical maximum downside levels of EV/sales given a company's growth and margin profile.

High-quality software companies with substantial net cash balances are in some ways the perfect area to average down, when/if prices go low enough. We did this heavily in 4Q22, and in our 3Q22 letter laid out our reasons for doing so; the same reasons apply today. These are mission-critical products that customers pay for on a recurring basis and are very difficult to turn off, where the main debate is how fast revenues grow rather than if they grow at all; they are not very susceptible to margin pressures, and in fact margins tend to expand when growth slows; and both fundamental analysis based on terminal margins and numerous PE takeouts have shown that there is a clear floor on medium-term valuations. We are watching this area closely and would love to engage much more.

We are sitting here with a still-high cash balance. Until our cash balance hits zero, further stock declines that generate mark-to-market losses would allow us to buy more and ultimately accrue more to our benefit than if stocks just went up from here. (And even when our cash hits zero, we can still reposition capital from lower-IRR positions to higher-IRR ones, as well as potentially take on a small degree of leverage if exceptional risk/rewards warrant it.)

As I wrote in our start-of-year letter, green markets are OK, but deep red markets are what get us really excited. Our eyes light up and our idea generation work spins up into overdrive, as declining prices open up entirely new opportunity sets previously closed to us due to valuation. It's time to get to work!

## **HashiCorp**

One company we've been averaging down in is HashiCorp. After following it for a year, we bought a small position at \$28 in early September, when at 6.8x calendar 2024 sales it was down meaningfully from summer highs and looked cheap relative to potential, though not excessively cheap enough to warrant a full position; but really loaded up in the low \$20s when the price reached and then fell below 5x sales. Today, at \$19, HashiCorp trades at 4.2x 2024 sales, a below-software-median multiple for a company with decidedly above-median long-term growth prospects. HashiCorp has a large cash balance relative to market cap (\$1.2 billion net cash vs. \$4.2 billion market cap), which, in addition to fulfilling the other criteria listed earlier, makes it a perfect candidate for doubling down as the fixed cash balance causes the enterprise value to swing dramatically relative to changes in the stock price. While this constrains the upside potential relative to that for a more levered company, more importantly it constrains downside as well.

### *Product*

HashiCorp develops software essential to managing cloud infrastructure – as such, it is directly exposed to the still-early transition to public cloud. HashiCorp releases all its products first as open source, which helps them gain wide adoption, then later tries to monetize with paid enterprise versions that offer enterprise-related features as well as support. HashiCorp's two main products are Vault and Terraform, both of which solved major customer problems, are loved by practitioners, and absolutely fit our criteria in terms of being “great products.”

Vault helps companies manage *secrets*, which can be thought of as machine-to-machine passwords or credentials. Secrets are needed to be accessed constantly as pieces of software call out to each other, e.g., microservices to microservices, to databases, to 3<sup>rd</sup> party APIs, etc. Prior to Vault, managing the complex web of secrets and rotating them often (i.e., like changing passwords) was a huge pain, and failure to do so led to security lapses. Vault provided an elegant, difficult-to-copy, and, importantly, cloud-neutral way to solve this problem, and as a result open-source Vault has become widely adopted.

Terraform helps companies use code to manage cloud infrastructure (“infrastructure as code”); to write standardized instructions for what the parameters of cloud infrastructure should look like, rather than configuring each piece of infrastructure manually by, say, clicking AWS settings. Just like with Vault, Terraform allows multi-cloud infrastructure management, which makes it superior to using each individual public cloud's infrastructure as code tool. In addition to Terraform being inherently better and easier-to-use than each cloud's built-in tool, as well as being a layer that can be used on top of multiple clouds, having cloud-neutral provisioning tooling allows customers to credibly threaten multi-cloud usage, which leads to less lock-in and thus better negotiating leverage vs. AWS or other providers. (The same logic applies to Vault; the 50-100:1 scale of each customer's spend on public cloud vs. on HashiCorp means that additional negotiating leverage can, in theory, make these products pay for themselves.) Even more so than Vault, Terraform open-source is widely adopted and has dominant market share.

### *Monetization*

Dominance of open-source Vault and Terraform as critical infrastructure tools is not in question. The debate for the investment revolves around HashiCorp's ability to convert companies using its tools for

free into becoming paid customers, for which HashiCorp charges a lot (for customers paying >\$100k, the average customer is paying ~\$600k, with 107 customers paying >\$1 million and 2 customers paying >\$10 million). Thus, the bulk of our work has been on understanding what gets customers over the finish line to becoming paid customers, and our conclusion is that these products are essential, with adoption being mainly an issue of cloud maturity. Our core thesis is that over the long term, as they reach greater cloud maturity, large enterprise companies will not run core infrastructure components unmanaged and unsupported. This fits into a broader trend we see across the infrastructure software landscape around a predictable shift at each company from initial in-house solutions built on top of open source to managed solutions over time; a phrase we often hear from customers is a realization that “why should we reinvent the wheel” when a better solution has already been built at a vendor who is specifically focused on that solution, and that often would actually come at lower total cost vs. in-house when the cost of engineers maintaining that in-house solution is considered.

This analysis of monetization is also made easier because of the market structure – Vault Enterprise is a very difficult-to-replicate product that has essentially no cloud-neutral competition (other than from in-house solutions at each public cloud, which have far fewer features and have the structural disadvantages described above regarding having to manage multiple solutions and suffering cloud lock-in), and Terraform Enterprise has only limited competition, from small vendors who build on top of open-source Terraform in order to offer competing managed solutions, but who have a hard sell with large customers vs. a trusted large vendor (HashiCorp) that developed and controls open-source Terraform.

The value proposition of paid Vault is better than the value proposition of paid Terraform, which is reflected in our estimate that Vault is ~50% of HashiCorp’s revenue today (with Terraform being 40%, and a third product Consul being 10%). First, Vault is a mission-critical runtime product – if your secrets management goes down, your infrastructure stops working – which makes proper management tools and enterprise support more valuable. Vault also falls under the cybersecurity bucket, which is a priority area of spend. Second, Vault’s monetization strategy was well planned-out from the start; HashiCorp from the beginning kept key Vault features enterprise-only, whereas in Terraform they released too many features to the open source. Essential Vault features such as advanced access management (only allowing certain employees access to certain secrets, which is a critical setup for companies to be able to show when they undergo security audits) and secrets replication (replicating secrets across multiple servers and datacenters) are only available in the paid enterprise version. For these reasons, for large companies who take security seriously, Vault Enterprise is kind of a no-brainer, and paid penetration should continue to increase steadily. The main reason why penetration isn’t already higher today is that many companies are still in the earlier stages of shifting to public cloud, with the earlier stages often marked by a free-for-all by various developer teams with limited controls, and in contrast to some other vendors of cloud infrastructure that are adopted earlier, it takes reaching a later phase of the maturity curve for customers to realize that they need to centralize secrets management, and that Vault Enterprise’s features are essential to managing secrets in a sustainable way over the long-term.

Because Terraform is non-runtime (if your Terraform goes down, your infrastructure will continue working for the time being, though you will be unable to make changes), paid Terraform is a harder sell. However, over time, the necessity of having management features and paid support for core infrastructure applies to Terraform as well. Unsupported Terraform open source is in essence run on individual engineers’ laptops, which is a very non-enterprise way to manage core infrastructure (though companies can as an initial step rig together their own systems to manage Terraform without paying for it, often not well). Terraform Enterprise brings access controls, compliance measures, policy enforcement, centralized management, and importantly, enterprise support (one of the customers we spoke with said that support alone was worth the price of his HashiCorp Terraform contract – running core infrastructure unsupported was too great of a risk to take). Though it will likely initially have a slower adoption rate than Vault Enterprise, when we look to the long term we see high penetration of Terraform Enterprise as well.

### *Investment case*

The investment opportunity is that the penetration rate of HashiCorp products in large enterprise companies is still low today. HashiCorp has as paid customers 36% of the Fortune 500 and just 23% of the Forbes Global 2000. In contrast to some other software management teams that put out such pie-in-the-sky estimates of TAM so as to be useless, HashiCorp management has laid out a well-reasoned long-term framework for revenue: \$2 billion (vs. \$583 million for calendar 2023, per our estimate). This is driven by a core penetration assumption that 50% of the top 4000 companies in the world will become HashiCorp customers (vs 851 >\$100k customers today), with an average revenue per large customer of \$800k (vs. ~\$600k today), with the remaining 20% of revenue driven by smaller customers. If our view is correct that every large company needs to at least seriously consider Vault Enterprise and Terraform Enterprise as they become more mature in using cloud, this seems a reasonable target that leaves further room for growth (through further increases in the number of customers and usage within each customer). Reaching \$2 billion in revenue would mean that HashiCorp currently trades at just 1.5x that target. For our exit valuation, we forecast a 2027 revenue of \$1.3 billion, a 22% CAGR from this year's levels, which would make the current EV just 2.3x that number.

The investment case isn't perfect; otherwise, it wouldn't be trading at 4x sales. Purchases of HashiCorp enterprise products, especially paid Terraform, fall on the more discretionary side because customers can jig together open-source versions that, while not ideal, can tide them over during times when budget is tight. As a result, versus some other enterprise software companies, HashiCorp has been disproportionately affected by this current environment, with growth slowing from near 50% in 2021 and 2022 to ~23% in 2023. Furthermore, our numerous conversations with former employees indicate that management has some areas of weakness, especially in sales/go-to-market/product marketing. This has hope of being fixed, in particular with the recent hiring of proven software leader and longtime HashiCorp board member Susan St. Ledger to lead global sales. Already, she has laid out a much-simplified sales message (condensing seven products into just two categories: Infrastructure Lifecycle Management and Security Lifecycle Management) that resonates with us.

This setup may (or is, judging by the movement in the share price) be turning other investors off. However, as we wrote in our previous letter, our intense focus on product means we look for "buts" – situations where people say "the product is great, *but*..." and pass. Sometimes the power of a great product outweighs the *but*, and we believe this is one of those cases. It is much harder to fix a bad product (we tend to pass immediately) than it is to fix bad sales. HashiCorp's product is *great*, which leads to customer adoption and gives the investment its potential.

Founded just 11 years ago, the HashiCorp founders have created a set of products (HashiCorp actually has seven widely used products; we focused on Vault and Terraform as the needle-movers, but some others have longer-term potential) that have reached ubiquitous adoption, with 250 million downloads in 2022. Co-founder Armon Dadgar, who former colleagues describe as a mythical engineer who by himself wrote key pieces of Vault's architecture over brief spans, is intensely involved and as CTO is driving the company's technical development. Paid Vault and Terraform have evolved into compelling products that we believe with high conviction will have much higher adoption in the future than they do today.

### **Conclusion**

We've had a very active few months, and HashiCorp is just one of several exciting positions we've started or re-entered. Beyond just numbers, I can feel the improvement in opportunity set viscerally – my screen is filled with tabs of companies I want to investigate, my excitement is high and my working hours

are inadvertently slipping later and later deep into the night. I trust that this period of increased activity primes our portfolio for higher returns, and if the selloff continues then even more so.

As always, I am immensely grateful for the opportunity to manage our capital and am hard at work trying to compound it at the highest rate that I can.

Yours,  
Tim Liu