October 31, 2019

Dear Partners,

During the third quarter of 2019, Lightsail Partners produced a gain of 0.8% net of fees. This result compares to a loss of 2.4% in the Russell 2000 index, and a gain of 1.7% in the S&P 500 index. For the year to date, the partnership has returned 23.0% net, compared to 14.2% for the Russell 2000 and 20.6% for the S&P 500.

	Lightsail Partners LP (net)	Russell 2000 Total Return Index	S&P 500 Total Return Index
3Q 2019	0.8%	-2.4%	1.7%
2019 YTD	23.0%	14.2%	20.6%
2018 (from April 16th inception)	10.0%	-12.1%	-4.3%
Cumulative return since inception	35.3%	0.3%	15.4%
Annualized return since inception	23.0%	0.2%	10.3%

Please refer to the disclosures at the end of this letter.

Since its inception on April 16th, 2018, Lightsail Partners has delivered an annualized gain of 23.0% net of fees. Over the same time period, the Russell 2000 index delivered an annualized gain of 0.2%, and the S&P 500 index an annualized gain of 10.3%.

Portfolio Update

Contrary to appearances, this was a very eventful quarter for the partnership. We have been hard at work: At the end of September we participated in a private investment in the public equity of a software company headquartered in a developed market abroad. We committed 18% of the partnership's capital to the investment. While evaluating this deal, we entered into a non-disclosure agreement with the company that granted us a level of access to its data, employees, partners, and customers that would have been difficult to attain as arm's-length public shareholders. Upon completion of our due diligence, we concluded that the deal represented one of the most attractive investments we have seen in recent years. We are very excited about this opportunity, and we discuss it in a little more detail at the end of this letter.

There was also significant activity elsewhere in the portfolio. One of our smaller positions, Pivotal Software, agreed to be acquired by VMware fewer than six weeks after we made our investment last quarter. (For a detailed look at our experience with Pivotal, see the next section.) The gains generated by the Pivotal acquisition were offset by a decline in the share price of our medical device holding, but as this company's stock has fallen, we have been increasing our investment, and it remains one of our largest holdings.

Perhaps the most important recent development is that, for practical purposes, we are now fully invested, with 92% of our capital in securities at quarter-end. Even better, the majority of our positions are trading at extremely attractive prices. In our opinion, right now the portfolio is the best positioned it has been since the inception of the partnership. To put it bluntly, we're short of capital. We could put new money to work immediately at what we judge to be very high prospective IRRs.

An Extended History of Pivotal Software

Last quarter, we mentioned a new investment we'd made in an unnamed enterprise software company. That company was Pivotal Software. In August, just six weeks after we purchased our Pivotal shares, the company announced an agreement to be acquired by VMware, a software company that has a long, tangled history with Pivotal, for a price 41% higher than our average cost.

There was plenty to like about Pivotal when we invested this past June. It was a sizeable business, with over $700 million in run-rate revenue. Over $500 million of this revenue came from recurring, high-margin software subscriptions, with the rest coming from a services business that was terrific at driving adoption and usage of Pivotal's software. The company's net expansion rate (a measure of how much money existing customers are spending with a company relative to the year before) was a mouthwatering 143%, one of the highest numbers we'd ever seen for an enterprise software company. And, while the company was losing money as it invested in its business, there was over $850 million of cash on its balance sheet, enough to finance years of future losses if necessary.

Despite these good points, we were able to acquire our shares in Pivotal at a very attractive price because the company's operations were decidedly *not* firing on all cylinders. While Pivotal was still growing at a rapid rate, almost all the growth was coming from additional sales to existing customers. New customer growth had nearly evaporated: The number of new customers that Pivotal was signing up had declined every quarter since its IPO in 2018.

How could a product so useful to Pivotal's existing customers be so irrelevant to new prospects? Investors feared that the competitive environment had changed for the worse, and that a new technology authored by Google, called Kubernetes, was making Pivotal's offering obsolete. To many observers, this year's fiscal Q1 results confirmed these fears, and Pivotal's stock price fell 41% the day they were announced. At that point, we had been researching the company for several months, and our work had convinced us that Kubernetes was not the death blow some feared it might be — in fact, we thought Pivotal was likely to see an eventual resurgence in growth. We bought as the share price crashed.

Pivotal's product is a software development platform: software for making software. The term of art is "PaaS," platform-as-a-service. If a software development organization is a factory for making applications, a PaaS is a group of manufacturing robots; they automate tasks that formerly required custom code, dramatically improving developer productivity. To developers, the ease of using Pivotal's platform is summed up in this haiku tweeted by a Pivotal employee:

> *Here is my source code*
> *Run it in the cloud for me*
> *I do not care how.*

Pivotal's platform consists of multiple components which perform several technical functions.[1] One of the most important, application execution, is handled by something called a container orchestration system. On the Pivotal PaaS, that system is called Diego. Kubernetes, the alternative that had many investors worried, serves the same function.[2] There are many technical differences between the two, and we would be happy to bore or regale you with them if you have an interest (just give us a call). What matters to the investment case is that Kubernetes had steadily gained popularity with developers since its release in 2015, and Pivotal's PaaS did not work with it.

Early last year, Pivotal at last responded to this growing issue, releasing a standalone Kubernetes-based product called PKS. But PKS was just a container orchestration system, not a full-featured PaaS, and it didn't integrate with Pivotal's existing PaaS. Consequently, PKS wasn't sufficient to stem Pivotal's declining rate of new customer wins. In June of this year, Pivotal attacked the issue more comprehensively. In the same Q1 report that included their distressing results, they also announced the upcoming debut of an overhauled version of the core PaaS — with an alpha release slated for July — that would offer users the option of Diego *or* Kubernetes for container orchestration. Management was taking a hard but important step.

The road from an alpha to a final release can be long. Would Pivotal succeed in transitioning to Kubernetes? Our research suggested that they would. Most important to us was the fact that Pivotal is not, at its heart, a container orchestration company; it is a PaaS company, with container orchestration as just one of the functions it handles. The various components of the platform are curated by Pivotal to best meet customers' needs, and those components change over time. In fact, *Diego itself was not the first container orchestration system within the Pivotal platform.* It was adopted in 2014 to replace an earlier system. This sort of product transition is rarely simple, but Pivotal had already proven it could change without too much trauma to the company. There was no reason to think it couldn't do so again.

Moreover, Pivotal's entire DNA is about continually improving software. It has built a business out of helping developers make better software, and it does this very well. In serious coder circles

---

[1] These include things like deployment automation, identity management, application routing, and logging. If you're interested in the details, you can start here.
[2] A good primer on containers and Kubernetes, written by Google, is here.

Pivotal is well known for the technical excellence of its people. The most powerful example of this technical prowess is the creation of Pivotal's software arm itself. When Pivotal was founded in 1989, it was a consulting company, teaching best practices for writing software via methodologies such as pair programming and agile development. In 2012, it was acquired by the technology conglomerate EMC, who also owned a controlling stake in VMware (yes, the same VMware that would eventually acquire Pivotal).

EMC took certain technology assets developed by VMware and handed them to Pivotal — assets that now make up the core of the Pivotal PaaS. But Pivotal's new assets were not exactly a windfall. They were weak and clunky programs, and had been passed to Pivotal only after VMware had tried, and failed, to rewrite them into usefulness. In a post on a technology discussion forum, one consultant to the company referred to the VMware technology as "poorly-designed abandonware." Pivotal, already experts in teaching others how to efficiently and effectively make good software, turned their expertise inward, and succeeded in transforming the suite into a powerful, modern PaaS — one robust enough to support major organizations like Allstate, Boeing, GE, and Wells Fargo, eventually generating half a billion dollars annually in subscription revenue. In the words of the same consultant:

> *[The] project should have failed. A huge, incredibly complicated body of enterprise software with near-100% team turnover? I would have bet against it ever working. But all that pair programming and rotation ... etc. just eventually ground the problem down. It ... took years but it looks like a success story now. I don't know of any other big takeover project like this that worked. It's a huge credit to the people working on it, and yes - to the "Pivotal process."*

Pivotal's lineage of self-transformation made it clear to us that the task of adopting Kubernetes was eminently achievable.

The other pillar of our investment thesis was that Pivotal's sales reach was about to expand through a deepening of its partnership with VMware. In 2015, Dell bought EMC, bringing Pivotal, VMware, and Dell into a single ecosystem. The three companies then began "jointly" marketing and selling their products. In practice, this meant that Pivotal's salespeople knocked on the doors of Dell and VMware's customers, without involvement from the sales teams at Dell or VMware. But this arrangement was changing for the better: as PaaSes became more important in the technology world, Dell and VMware began getting more involved with Pivotal. Crucially, VMware's massive sales force would now begin selling Pivotal products. Success here would take time, but would open the door to a much deeper relationship between the two companies.

A key sign that Pivotal was becoming a priority for the Dell complex of companies came when we spoke to a CIO who was already a VMware customer. This CIO had just met with Michael Dell, and told us Mr. Dell had made it a point during their meeting to pitch Pivotal specifically. Pivotal and its products now had the attention of the CEO of Dell itself. We believed that once the combined Pivotal/VMware/Dell sales teams had a full PaaS employing Kubernetes to sell, there was potential for an acceleration of new customer wins at Pivotal.

Nevertheless, Pivotal's share price implied that it would never recover from its Kubernetes-induced slump. We saw opportunity. At our purchase price, the company had an enterprise value of $2.2 billion, representing a multiple of just 5x the gross profit it generated in the trailing twelve month period. To put this fire-sale valuation in context, the most comparable recent transaction — IBM's acquisition of Red Hat in late 2018 — happened at 12x gross profit.

Sure, Pivotal deserved some discount: a key component of its main product was undergoing a major overhaul, and its controlling owner, Dell, had a history of using rough elbows when dealing with minority shareholders. But a 58% discount to Red Hat's private market valuation seemed so disproportionate to us as to border on egregious.

We made Pivotal an 8% position at cost. Six weeks later, VMware made their acquisition offer — likely seeing what we saw, albeit at much closer range — and as a result we did not own the shares for long. Pivotal would only report one more quarter as a public company, so it was too soon to tell whether our thesis about Pivotal's ability to adopt Kubernetes was correct, but in the months since our purchase, we watched with interest as the company's new Kubernetes product went into alpha. When Pivotal announced its fiscal Q2 results in September, the company reported its first increase in new customer wins since going public.

We would have preferred to see Pivotal's performance prove us right over the long term, and to participate in the company's growth as it reaccelerated. Instead, we recorded a nice profit in a few weeks, and then turned to find a new investment into which to redeploy our cash.

Lightsail's First Private Investment in Public Equity

Late in September, we participated in a private financing deal for a small publicly-traded enterprise software company located in a developed market outside of the United States. The company is currently sub-scale, but has a large growth opportunity ahead of it that will require additional capital to pursue, hence the capital raise.

Despite the company organically growing at 30-40%, generating gross margins above 70%, and demonstrating high customer retention, we were able to make our investment at a price that valued the company at less than 3x its trailing revenues, a multiple virtually unheard of for subscription software companies here in the United States. We believe this striking gap exists because the company is domiciled in a country that has not yet produced many big tech companies. As a result, that country's stock market is lacking in sophisticated software investors, especially at the low end of the market capitalization spectrum. If this company were listed on a U.S. stock exchange, we are certain the deal would have been completed at a *much* higher price than what we paid; the company's closest peer was recently acquired for 5.5x its revenues.

This company sells subscription software that integrates important IT and telephony capabilities into its customers' CRM systems. Because we'd love to own more of the company should the opportunity arise, we'll save the identifying details of what they do for a later discussion. We'll summarize by saying that when we first encountered this company, we had an initial skepticism that their product might be easily replicable, and that it might be helpful, but not essential, to the customer. As we spoke to the company's management, its customers, its major CRM partner, and its competitors, we quickly became disabused of these concerns.

Our company's product offers a superior daily workflow, bringing unquestionable value to the customer. Moreover, any attempt by another company at competitive displacement would be an uphill battle. An upstart looking to challenge our company would first need to write a good piece of software that integrated with numerous CRMs, while simultaneously assembling a reliable global IT and telephony infrastructure. Then it would would have to gather capital to finance years of ongoing losses as it developed the product and built a reputation in the market. Once the product was finally ready for business users, the upstart would have to convince a large swath of customers to move away from the known and respected market leader (our company). And the upstart would need to do *all* of the preceding well enough to convince major CRM partners to deemphasize their relationships with our company in favor of the upstart. We think it is highly unlikely that a new entrant would do most of these, let alone all. Anyone looking to gain a leading position in this market would quickly discover that the easier, faster, and more prudent way to get there would simply be acquiring our company at a fair price.

Our base assumption is that in four years, the intrinsic value of this company's shares will be 3x what we paid. We think we are very unlikely to lose money given the stark undervaluation of the stock, and the stability of the underlying business. However, the company's stock is highly illiquid, with meaningful positions currently attainable only through negotiated block transactions. We are prepared to weather a high degree of share price volatility until the company establishes a larger base of software-inclined investors. To the extent that this volatility occurs when we have excess capital, we will welcome it.

Portfolio Composition

At the end of September, we had invested 92% of Lightsail's capital in securities, with the rest held in cash. Our two largest positions (our cable investment and the new software investment discussed above) were each 18% of our portfolio, and our top five positions represented 72% of the portfolio. In addition to our new investment in the enterprise software company, we also increased our holding in three existing investments during the quarter.

We now have eight notable holdings: two broadband providers, two enterprise software companies, one patent-protected medical device company, one provider of share registry services, one building products manufacturer, and one specialist bank. At quarter's end, the companies in our portfolio had a weighted average market capitalization of $1.1 billion.

We look forward to reporting to you on our partnership's progress again early next year.

Sincerely,

James Basili & Torin Eastburn
Managing Partners

Disclaimer

Lightsail Capital Management LLC (the "Investment Manager") is a private fund advisor exempt from registration. The Investment Manager provides discretionary investment advisory services to Lightsail Capital Partners LLC (the "General Partner") and, through the General Partner, to Lightsail Partners LP (the "Partnership"). Lightsail Capital Management LLC and its affiliated entities ("Lightsail") are providing this document on a confidential basis, and for informational purposes only. Any reproduction or other distribution of this material in whole or in part without the prior written consent of the Investment Manager is prohibited.

This document shall not constitute an offer to sell or the solicitation of any offer to buy any security or investment product. Any such offer or solicitation to invest in the Partnership may only be made by means of delivery of definitive offering documentation, which contains important information (including investment objective, policies, risk factors, fees, tax implications and relevant qualifications), and only in those jurisdictions where permitted by law. In the case of any inconsistency between the descriptions or terms in this document and the definitive offering documentation, the offering documentation shall control. These securities shall not be offered or sold in any jurisdiction in which such offer, solicitation or sale would be unlawful until the requirements of the laws of such jurisdiction have been satisfied. Due to, among other things, the volatile nature of the markets, an investment in the Partnership may only be suitable for certain investors. Parties should independently investigate any investment strategy or manager, and should consult with qualified investment, legal and tax professionals before making any investment.

An investment in the Partnership is speculative and involves a high degree of risk. Opportunities for withdrawal/redemption and transferability of interests are restricted, so investors may not have access to capital when it is needed. There can be no assurance that the Partnership's investment objectives will be achieved, or that it will or is likely to achieve results comparable to those shown herein, or will make any profit, or will be able to avoid incurring losses. Past results are no guarantee of future results. All investments involve risk, including but not limited to the loss of principal. Investors must be prepared to bear the risk of a loss of their investment.

The description of investment strategies in this document is intended to be a summary and should not be considered an exhaustive and complete description of the potential investment strategies used by Lightsail. Specific companies or securities shown in this document are meant to demonstrate Lightsail's investment style and the types of situations and instruments in which we invest and are not selected based on past performance. Positions reflected do not represent all the positions held, purchased, or sold by the Partnership, and in the aggregate, the information may represent a small percentage of activity. There is no assurance that the Partnership will make any investments with the same or similar characteristics as any investments presented. The investments presented are not a reliable indicator of the future performance or investment profile of the Partnership.

While Lightsail has used its best efforts in preparing this document, and believes all the information herein to be accurate and from reliable sources, it makes no representation or warranty with respect to its completeness or accuracy, nor can Lightsail accept responsibility for errors appearing in the document related to information obtained from third parties. Any projections, forecasts, or estimates in this document are necessarily speculative in nature and are based upon internal analysis and certain assumptions which reflect the views of Lightsail as of the date indicated, and should not be construed to be indicative of actual events that will occur. No representations or warranties are made as to the accuracy of such forward-looking statements. Any statements of opinion constitute only current opinions of Lightsail, which are subject to change and which Lightsail does not undertake to update.

Disclosure

The performance information herein has been prepared by or on behalf of Lightsail and has not been independently audited. All performance results for the Partnership are calculated based on the aggregate net asset value of fee-paying investors only, and are presented net of management fees charged, brokerage commissions, administrative expenses, and accrued performance allocation, if any, and include the reinvestment of all dividends, interest, and capital gains. While performance allocations, if any, are accrued monthly, they are deducted from investor balances only annually or upon withdrawal. The performance results represent fund-level returns, and are not an estimate of any specific investor's actual performance, which may be materially different from such performance depending on various factors, including the timing of contributions and withdrawals and any special terms granted to an investor. All performance results are estimates and are subject to future adjustment and revision. Investors are encouraged to review their capital account statements to obtain their individual returns. The Partnership's inception was on April 16, 2018.

While the performance of the Partnership has been compared here with the performance of well-known and widely recognized indices, any index or benchmark comparisons herein are provided for informational purposes only and should not be used as the basis for making an investment decision. There are significant differences between Lightsail's strategy and the benchmarks referenced, including, but not limited to, risk profile, liquidity, volatility and asset composition. The index performance results shown include the reinvestment of dividends.