

# Pure's Innovation Plan: Delivering Data Advantage



When we started Pure in 2009, we laid-out a bold but simple plan. We saw the coming of NAND flash technology as an opportunity to re-think storage entirely, so we used flash as an opportunity to build a new class of [smart data storage](#) solutions, focusing our innovation in three key areas:

1. **Build affordable all-flash storage.** We pioneered storage software that, when coupled with flash, could deliver all-flash storage capable of delivering a 10X improvement to cloud-era datacenters (performance and efficiency) at a cost profile that enabled mainstream adoption to replace both Tier 1 and Tier 2 storage.
1. **Make data storage simple and “self-driving.”** By first building a fundamentally simpler, automatable, and always-on storage platform, and then combining it with SaaS-based management, orchestration, analytics, and predictive support technologies - we set out to change the storage ownership experience, making storage “just work” - whether you are a virtualization admin managing just one array or are a storage expert automating 100s of arrays in a large-scale cloud.
2. **Make data storage “Evergreen”™.** Since flash was changing so fast, we also realized that we had to eliminate the 3-5 year “forklift upgrade” cycle that dominated the storage business model, so we built a software-driven storage architecture that was built on an agile hardware approach to be continuously upgradable, and coupled it with innovative business models to make those upgrades accessible to customers.

Seven years later, our original Pure mission is well on the path to being accomplished. We're proud to have had an outsized impact on the \$35B storage market, as the storage behemoths have all pivoted to chase Pure - attempting to emulate both our product strategy and business model. **It's time to push the industry forward again - with a new set of innovation priorities.**

As we look at the next chapter for Pure and for the industry, we're focused on a central theme: *exponentially expanding the value of data*. We aim to deliver a data platform that enables our customers to build a new class of applications and to extract new insights from data - we call this *data advantage*. To do this we're going to need to expand our reach to new types of information, to deepen our integrations with next generation applications and platforms, and to tailor our solutions to the developers of tomorrow's clouds and applications. With *data advantage* as our focus, we've set **four (additional) new innovation priorities for the next 5 years at Pure:**

1. **Bring flash to large-scale unstructured data.**
2. **Transform application data storage (again) with NVMe.**
3. **Deliver next-generation converged infrastructure, optimized for cloud and "new stack" applications.**
4. **Make application development and data protection effortless in the cloud era.**

So why five years and not the next decade? Well, no one can see that far ahead, and we still pride ourselves with innovating at "startup speed." Keep reading to dive deeper into each of these four new innovation investments. And to be clear - we're far from done with our first three innovations areas! The storage industry is still only about 15% into the all-flash transition, and we're continuing to innovate with SaaS-based management, predictive analytics, and even more business model transformation.

## **But first...is storage even strategic?**

As the businesses of legacy storage companies continue to falter, some have asked 'why build a storage company at all?' Our belief is simple: in the increasingly digital economy *data is key*. Independent of whether your infrastructure is public cloud, private cloud, SaaS, or most likely some hybrid of all three - *data is key*. Independent of whether you are doing business online, delivering better patient care, sequencing the next discovery, delivering your software via SaaS, or evolving your business analytics to real time - *data is key*. Data is the strategic core that enables competitiveness and differentiation in tomorrow's economy; *data is where the potential lies*. Yet we are just at the beginning of leveraging machine learning and advanced analytics to take data insights beyond even human prediction. Differentiation lies in having the most data, analyzing it the fastest, discovering new insights, and ultimately delivering new products and experiences not possible without it.

Yet in most data centers today, data is also the key bottleneck. Data (on disk) is the last mechanical gate in the compute chain. Data is where compromises are ultimately made: how much to store, how much can be analyzed, how long to wait for answers - compromises no one wants to (or should have to) make.

Today's IT landscape is dominated by vendors who are treating data as a compromise. Legacy storage vendors are retrofitting flash to decades-old systems, and perpetuating the status quo vs. embracing the true potential of flash. IT behemoths are focusing on becoming "vertically integrated" to sell customers one of everything vs. being best-of-breed at anything. Public cloud vendors are making it cheap to store data but expensive to use it. And hyper-converged vendors are focusing on reducing cost by making storage just an integrated commodity function vs. a platform for strategic advantage. These vendors are all pivoting to low cost/low differentiation data strategies, vs. investing in a platform to drive *data advantage*.

Pure believes that there's a place in the market for a pure-play data innovator, wholly-focused on driving *data advantage*. Let's explore the investments that will help us get there.

## **Innovation Area 4: Bring flash to large-scale unstructured data.**

To date, Pure's success has been built upon helping core business applications and the clouds that run them be faster and more efficient. But this mission ignored the majority of data - the vast pools of unstructured data that in today's world of digital sensors and machine-generated everything is being created faster than it can be analyzed.

Earlier this year we introduced FlashBlade™ - the second major product from Pure Storage - aimed squarely at tackling the "big data" problems and opportunities of tomorrow. We've shared our thinking extensively about why this investment is such a strategic bet for Pure, why we think the market is as big or bigger than the current market for FlashArray, and why we believe we've developed a technical architecture that is unsurpassed in this space.

Since first announcing FlashBlade at Pure//Accelerate 2016 in March, we've shipped the Directed Availability version of the product to production deployments, and have already seen some truly mind-blowing results with customers who are pushing the boundaries of machine learning, genomics, large-scale data analytics, software development, and creativity.

FlashBlade development continues aggressively, and we are still anticipating releasing the GA version this year as announced previously. Give us a call if you'd like to try one today!

## **Innovation Area 5: Transform application data storage (again) with NVMe.**

Our founding vision for Pure was to eradicate spinning disk from the data center, but it turns out that the remnants of disk remain in the storage architecture even after the disk is gone. Legacy protocols, most notably SCSI, remain the archaic way we still "speak" to both flash within storage arrays, and to storage arrays over the storage network. It's time to banish SCSI for good.

The advent of the new NVMe and NVMe/F protocols hold the promise to do just that. NVMe (over PCIe) enables storage arrays to eliminate SAS (serial-attached SCSI) links from their controller processors to flash, and NVMe/F (NVMe channeled over RDMA-enabled fast Ethernet) enable dramatically faster, SCSI-less connections between servers and storage arrays, making them "appear local" to servers and enabling the best of both worlds: local storage performance with networked storage shared efficiency and manageability.

NVMe and NVMe/F have huge potential to open-up the next generation of advances in storage-centric computing, but they face the same challenges that flash did a decade ago: cost, complexity, and interoperability. Early attempts at NVMe are once again focused at the ultra-high-end of the market, limiting its market potential.

Pure's approach to NVMe is simple, and perhaps unsurprisingly similar to how we first approached flash:

- **We believe NVMe is for everyone.** We're building a set of NVMe-enabled products that will

allow mainstream NVMe adoption across a wide set of use cases – not just for the performance elite.

- **We believe NVMe shouldn't be a compromise.** Early all-flash arrays were all performance and no features. Just like Pure pioneered the feature-rich AFA enabling the technology to mainstream – we believe that NVMe solutions should incorporate all the critical data management features that customers need and expect (HA, RAID, data reduction, encryption, snapshots, and replication).
- **We believe NVMe should be an Evergreen evolution.** Others will create “new” generation products for NVMe, leaving their existing customer's investments behind. The move to NVMe will be a multi-year transition, and we're committed to taking our customers there in a non-disruptive, investment-protected, Evergreen fashion.
- **We believe NVMe will grow the market for shared storage.** Many of the “new stack” database and analytics applications developed in the past decade have taken a DAS-centric architectural approach, and as a result often were deployed on whitebox servers with DAS flash cards or SSDs vs. shared storage. NVMe/F + NVMe enable these applications to leverage what looks like server-local storage, but to now take advantage of shared storage resiliency, efficiency, and manageability.

You'll hear a lot more from us in 2017 about NVMe. We anticipated the NVMe transition 3+ years ago when we first designed FlashArray//M, and architected it for the non-disruptive evolution to NVMe from the start, and we recently put our money where our mouth is and issued the industry's first NVMe-Ready guarantee. Put simply – given the coming rapid transition to NVMe, we don't believe anyone should buy a storage array in 2016 that isn't ready to upgrade to NVMe in 2017, and we're guaranteeing that FlashArray//M is NVMe-Ready today. And if you'd like to dive technically deeper on NVMe – tune in next week to [our second episode of Outside the Box](#) (our innovation show by techies for techies).

## **Innovation Area 6: Deliver next-generation converged infrastructure solutions, optimized for cloud and “new stack” applications.**

Storage doesn't exist as an island – it's deployed with compute, networking, virtualization, and all this is ultimately in the service of applications and users. Nearly a decade ago the concept of “converged infrastructure” was born – an attempt to enable customers to continue to adopt best-of-breed technologies at every layer, but to gain the advantages of an end-to-end tested stack.

Since the advent of the first generation of converged infrastructure, cloud computing architectures were born, virtualization became pervasive, and flash transformed the storage architecture. Applications also evolved, and a new class of database and analytic applications have been developed, many of which simplify legacy heavyweight relational databases and build-in “light” storage management themselves. We see an opportunity to re-think CI solutions for this modern cloud era, an era that demands greater agility and uses public cloud as the key comparison yardstick.

It was this need to modernize that first inspired us to create FlashStack Converged Infrastructure Solutions over two years ago, and Pure and Cisco Systems recently announced an [expansion of our collaboration](#) around FlashStack, new solutions for databases, virtualization, and VDI, as well as significant joint customer momentum. While FlashStack has been very successful to date, it continues to evolve to support new use cases, in particular focusing on private cloud – we believe FlashStack is the most reliable, powerful, and

agile private cloud foundation available. FlashStack is a major joint innovation opportunity, and we plan to continue to drive innovation across the CI solutions market in the coming years as we define the next generation of CI.

But looking even further beyond, we see an opportunity to re-define the lines between databases, analytic stores, and data storage. In the end – they are all about interacting with your data – shouldn't they be much more seamless? Can we blur the lines to create a new “data tier” to power next-generation applications? Stay tuned – we'll share more of our thinking soon.

## **Innovation Area 7: Make application development and data protection effortless in the cloud era.**

Delivering reliable, performant, and efficient storage is only a very small piece of the data lifecycle. It doesn't matter how easy our storage is, if customers then have to go through painful hoops to make it accessible to application developers at the beginning of the data lifecycle, and then manage copies and copies and copies of stranded data to check the box on backup, DR, and compliance at the end of the lifecycle. While storage has improved, these data management processes have remained largely unchanged for decades, and drive much of the proliferation of unnecessary data copies in the enterprise. It's also a whole new era – a cloud era – where customers are increasingly looking to leverage the public cloud to both develop applications and to retain data for long-term protection. So we ask ourselves a set of “what if” questions:

- What if application developers could have access to unlimited, instant, and space-efficient copies of their data? What if they could control and manage these copies without burdening a storage administrator?
- What if these data copies could be automated as part of a DevOps cycle? What if your storage array (and your next-gen CI) was just another piece of infrastructure that acted as code in your environment?
- What if your on-prem private cloud infrastructure could be seamlessly integrated with the public cloud? What if data could be easily moved between on-prem and the cloud, enabling development anywhere, deployment anywhere, and burst capacity anywhere?
- What if this data movement could be orchestrated end-to-end? Delivered as a [SaaS](#)-managed experience, with custom-tailored views for everyone involved in the process, enabling self-service and agility?
- And what if all this was just as easy as Pure is today, and just worked?

We see the opportunity to link our on-prem storage arrays seamlessly with the public cloud, integrating application development, DevOps, and data protection workflows seamlessly with the storage. This enables speed and agility in application development, and allows customers to shift their workloads between private and public clouds as it suits their needs. Stay tuned – we're working on this and collaborating closely with like minded best-of-breed partners on the development and delivery of a modern and re-imagined solution. One thing you can bank on is that will make it all really, really, really easy.

## **Creating Tomorrow's Pure-Play Data Innovator**

So what does it all add up to? We believe that the potential of data is so huge, so integral to the future of the digital economy, that the market deserves a pure-play (pun intended) company that focuses wholly on delivering *data advantage*. Big data. Fast data. Data analytics. Data science and machine learning. Data

orchestration between clouds, and data lifecycle management. We are Pure, and we *deliver data advantage*.