

Black Duck Software helps more than 2,000 companies worldwide secure and manage open source software. To do so, Black Duck relies on the data in its KnowledgeBase™, the world's most comprehensive database of open source license and vulnerability information. Black Duck accesses and updates the KnowledgeBase 24/7, making the storage platform for all that data a vital element of the company's IT infrastructure. Black Duck has standardized on FlashArray™ and FlashBlade™ products from Pure Storage® for their performance, non-disruptive scalability, effortless management and cost-effectiveness.

BLACKDUCK

BUSINESS TRANSFORMATION

By deploying FlashBlade and FlashArray systems, Black Duck can introduce new services for customers more quickly, cut capital and operating costs, and reduce the burden of storage management.

GEO

North America

INDUSTRY

Software

“FlashBlade met all our criteria — high performance, a reduced data-center footprint, very low power draw, and an architecture optimized for unstructured data.”

Kevin Smith, *Senior Director of IT*

BLACK DUCK SOFTWARE SELECTS PURE STORAGE TO REDUCE STORAGE COMPLEXITY

Open source software use is exploding worldwide because of its economic, productivity and time-to-market value. But the benefits of open source are accompanied by significant security and management challenges that include meeting license requirements and identifying open source vulnerabilities. That's why organizations worldwide rely on Black Duck Software to automate the processes of securing and managing their open source.

The heartbeat of Black Duck's operations is the KnowledgeBase, a data warehouse of more than 120TB that contains virtually every version of more than 2 million open source projects. “We use a wide variety of electronic and human resources to continuously add to and update the KB. That includes a team of ‘spiders’ who crawl the Web 24/7 looking for open-source repositories, and then download them into the KnowledgeBase,” said Kevin Smith, Senior Director of IT.

Black Duck's flagship product is The Hub, which leverages the KnowledgeBase data to provides an inventory (Bill of Materials) of the open source code in an application or container. It also identifies any known open source vulnerabilities and license compliance challenges. In addition, it provides remediation guidance and dynamically monitors the inventory for any new vulnerabilities that affect the inventory code.

The KnowledgeBase is accessed tens of thousands of times each day, as virtually all Black Duck's products rely on it. Given the criticality of the KnowledgeBase to the company and its customers, the storage infrastructure on which it resides must operate with speed, reliability and high availability.

HIGHER-PERFORMING STORAGE TO MEET INCREASING DEMANDS

Black Duck recognized that its legacy storage infrastructure needed updating. “Because the KB grows continually, we always need adequate headroom in our storage system,” said Smith. “That was hard to achieve with traditional mass storage, because it was time-consuming and disruptive to add capacity. Plus, legacy business models weren't friendly to incremental increases in capacity over time. We needed something different.”

The search for an alternative approach led to Pure Storage. “We decided to try a FlashArray in our development environment to get an idea of how it would work,” said Smith. “We saw great results.”

COMPANY:

Black Duck Software
www.blackducksoftware.com

USE CASE:

- Database – Postgres®
- VSI – VMware® vSphere®

CHALLENGES:

- Ever-increasing data volumes required more storage, but adding capacity proved to be difficult, complicating long-range IT planning.
- Insufficient IOPS on spinning-disk storage system limited utilization of data resources.

IT TRANSFORMATION:

- Faster data imports to the KnowledgeBase provide DevOps customers with up-to-date information on millions of open source projects.
- Data-reduction rates as high as 19:1 reduce the footprint and cost of storage.

“FlashBlade has no problem handling the concurrent traffic.”

Kevin Smith, *Senior Director of IT*

In addition to significantly lower latencies, the Pure Storage array allowed Black Duck to move forward on a strategic IT goal — virtualization. “Originally, most of our compute environment was physical because of the demands of our systems,” Smith noted. “By putting in the Pure Storage FlashArray, we have eliminated the need for physical servers, and have been able to reach almost 90% virtualization on our servers,” using VMware vSphere.

Within six months of its initial installation, the FlashArray //M series model was expanded twice — both times with no production interruption — and a second array was purchased for the production data center. “The seamless upgrades with Pure Storage have been a real benefit,” Smith said. “We do it live in just a few hours, as opposed to taking weeks or months, which is typical with traditional storage. Anytime we’ve had to increase capacity, it’s been invisible to end-users.”

FLASH ARRAYS LOWER BOTH CAPITAL AND OPERATING COSTS

A significant benefit of the Pure Storage arrays are high data-reduction rates. Across all applications and databases, Black Duck has seen data-reduction ratios of around 9:1, which Smith said, “blows everything else we ever saw or tested out of the water.” He added, “For our replicas of Postgres on the FlashArray — which are effectively live backups — we are seeing 19.6:1 data reduction.”

The high data reduction means savings on both capital and operating expenditures. “If we weren’t getting these great reduction rates, we’d have to buy three to four times the storage we currently have,” Smith noted. “Installing the Flash Array allowed us to shrink storage from three racks down to just 10U. As we use a co-location facility, that means significant savings on rack space, power and cooling costs.”

Another factor adding to the cost-effectiveness of a Pure Storage solution is the Pure Evergreen™ Storage program. “With Evergreen, we won’t have to go through arduous forklifts to replace anything, and we’re assured of always having the latest technology from Pure,” Smith said. “And it’s all done without downtime while preserving our investment. It’s a complete departure from traditional business models.”

FLASHBLADE ENABLES CONTINUING GROWTH

With continuing growth in both the number of customers and the size of the KnowledgeBase, Black Duck needed to determine how it would meet its storage needs in the future. “When the time came to decide our next steps for storage, Pure Storage just made sense, based on our previous experience,” Smith said. “Our top criterion was the number of IOPS. Next, was the ability to easily expand the system at the rate we need to.”

Pure Storage presented Black Duck with details about what at the time was a new product, FlashBlade™. FlashBlade is an elastic scale-out system that delivers all-flash performance to multi-petabyte-scale data sets. It combines commodity hardware components with unique software and ultra-dense packaging to achieve extreme performance, scale, density and efficiency.

“The more we learned about FlashBlade, the more we liked what we heard,” Smith recalled. “It met all our criteria — high performance, a reduced data-center footprint, very low power draw, and an architecture optimized for unstructured data, of which we have a lot. And, it supports NFS (Network File System).”

Open-source code is now downloaded to servers through FlashBlade, whose high bandwidth and IOPS make it ideal for processing very large amounts of unstructured data. The FlashBlade system then feeds the data into a FlashArray, whose architecture is

“With Pure Storage, you know they will handle it. It gives you peace of mind.”

Kevin Smith, *Senior Director of IT*

optimized for database transactions, such as those handled by The Hub that matches a customer's software assets against the KnowledgeBase.

Smith said Black Duck has found compelling advantages to using both FlashArray and FlashBlade. “The FlashArray is great for transaction-oriented database workloads, and has allowed us to virtualize our server environment, all while greatly reducing rack space,” Smith said. “FlashBlade is an excellent NFS solution that addresses our need to support concurrent workloads with high IOPS requirements. Plus, it allows us to expand capacity easily. Best of all, both systems can be managed from the same easy-to-use interface.”

On an operating basis, FlashBlade has delivered the high performance required by Black Duck. “We really need high IOPS, because we not only have lots of open-source code being downloaded into the KnowledgeBase, we also have customer code hitting it all day. On an average day, we can have hundreds of simultaneous calls to the KnowledgeBase, and FlashBlade has no problem handling the concurrent traffic.”

Simplified storage expansion has also been a key benefit to Black Duck. “In a matter of days, I can get a quote from Pure Storage, place an order, and have new storage installed within days — all without downtime,” Smith said.

EFFORTLESS MANAGEMENT AND SUPERIOR SUPPORT

Another significant benefit of Pure Storage is its effortless management. “On a daily basis, we don't do anything with the Pure arrays except look at the informative graphics,” Smith said. “Most of the team has the Pure1® app on their phones, which means anyone can check the status of the arrays at any time. But we hardly even do that, because Pure is looking at our systems 24x7, and will make us aware of an issue even before we know it.”

One such incident occurred when one of the FlashArrays hit 108% of capacity. “Pure Storage customer support initiated a call to one of our engineers to explain that the array was capable of handling this load level without interruption or loss of data, although obviously it isn't recommended,” Smith reported. “They got it down below 100% in very little time and added capacity soon after that. So, what would have been a major crisis with any other vendor turned into a minor fix.

“With Pure Storage, you know they will handle it. It gives you peace of mind.”

Smith counts numerous benefits of using Pure Storage. “The cost per terabyte is amazingly attractive. The operating costs are far lower. I don't need to dedicate a resource for SAN management. And, I can add capacity whenever we need it without a minute of disruption. I've been in IT for many years, and I don't know that I have ever had the same level of vendor loyalty as I do toward Pure.”



info@purestorage.com
www.purestorage.com/customers