panasas®

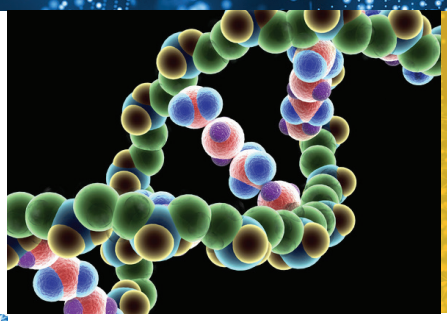

BUYER'S GUIDE | Life Sciences

Breathe New Life into Storage Infrastructure

Selecting Scale-out NAS Solutions for Life Sciences R&D Organizations

This buyer's guide is intended to help you specify, evaluate, and choose your next scale-out network-attached storage (NAS) system. First, let's look at why you should consider upgrading your storage infrastructure in terms of business benefits to your life sciences organization.

How New Storage Technologies Can Drive Research Success

The volume of life sciences data is growing at a phenomenal rate, driven by disciplines such as genomics, predictive biology, drug discovery, bioinformatics, computational chemistry, clinical trials, and personalized medicine. By 2019, 30 percent of the top 100 life sciences R&D organizations will have big data projects in production.[1] Research organizations that want to lead their peers into the future will make plans now to avoid being overwhelmed by the enormous potential of genomic data. Here are the potential benefits of upgrading your storage infrastructure now.

Stay One Step Ahead of Your Researchers' Storage Needs

The first human genome took 10 years to decode. Today, the same data can be generated in just three days. As the rate of information acquisition rapidly accelerates, so do the demands of your researchers for ever-shorter cycle times from data to results. Researchers want to publish their results now. Building the right storage foundation positions your organization for groundbreaking scientific research insights using modern rapid-iteration discovery methodologies.

Create Differentiation by Using Both Public and Private Data

The rising volume of global health data—about 50 percent growth per year by one estimate—is being fueled in part by public sources of scientific information. Companies that fail to take advantage of this open source information risk losing market share to their more aggressive competitors.[2] Your organization's success in this area depends in part on having a powerful and scalable storage infrastructure that makes it easy to stay ahead of your data growth needs.

Maximize the Productivity of Your Research Staff

Having the right information is necessary for good research, but it's not enough. Your researchers must be able to perform sophisticated processing-intensive analyses to extract the scientific insights hidden within the data. Legacy storage infrastructures can create a significant bottleneck that chokes performance of compute-intensive life sciences applications and limits the productivity of your most valuable resource: your top scientists' time. By connecting genetic and bioinformatics data directly to your computing grid instead of copying from the storage infrastructure to the compute grid, you can gain up to 50 times improvement in time to results. That's how you help your research teams publish results faster, leading to life-changing discoveries in less time.

How to Prepare Your Requirements Checklist

Most life sciences project managers looking for a storage solution organize their thinking in a checklist that guides both the architecture redesign and the vendor selection process. You can begin preparing your checklist by answering these four key questions.

1. How Will New Storage Technologies Impact Our System Architecture?

Start with an architecture that offers the greatest degree of flexibility for deploying applications, servers, storage, networks, and security. Then decide how your organization will scale in response to the need for additional performance and capacity. Consider how you will take advantage of private and public cloud resources to supplement your on-premise capabilities. Don't ignore storage system management, which has a considerable effect on increasing your operation expenses.

2. How Fast Will Our Data Grow?

Analyze historical patterns to characterize the general trend of your organization's data growth. Give significant weight to the data-generating potential of life sciences sequencers and applications that your organization plans to acquire in the next three to five years. For example, genomic sequencers from Illumina, PacBio, and others generate data 50 times faster than the previous generation of instruments of just a few years ago. Without the ability to scale storage quickly and gracefully, even one instrument acquisition could overwhelm your storage infrastructure.

"Meeting our commitments to researchers requires extremely high computational power that is available 24/7. Panasas lives up to its promise of terrific performance with negligible maintenance and administration time."

Dr. Warren Kaplan
Chief of Informatics
Garvan

---

3. How Much Performance Do Our Researchers Need?

Today, more than 25 percent of life scientists require high-performance computing (HPC) capabilities, and that number is growing steadily.[3] No matter how well your infrastructure performs today, it will not be enough in the near future. While the primary driver of performance is processor power, compute grids depend on the storage system to supply large data sets at the necessary rate. A legacy storage infrastructure not built for large-scale Linux computing can be a significant performance throttle for compute-intensive applications such as genomic sequencing and high-resolution medical imaging.

4. What Is Our Optimum Refresh Cycle?

As market research firm Forrester points out, "Up to 79 percent of organizations refresh their wired networking infrastructure every five years, guided by industry averages that originate from vendors."[4] Study your own ability to publish research data to determine the return on investment (ROI) and total cost of ownership (TCO) history and projected growth. This exercise can help identify the refresh interval that best fits with your specific storage experience and projections. By extending the life of storage investments, you improve overall ROI and free up resources to invest in other areas such as high-performance Linux compute farms needed to drive scientific breakthroughs.

How Panasas Stacks Up

The features and benefits of the Panasas® ActiveStor® scale-out NAS solution go far beyond those of other NAS offerings. As life sciences and bioinformatics data requirements expand, the stakes involved in choosing a storage vendor become higher and higher. Here are four key reasons why you should consider the ActiveStor solution as part of your strategy to analyze and publish more research data.

1. True Linear Scalability

With the Panasas solution, you can easily and seamlessly scale performance and capacity without limiting the number of compute clients on the network. Your budget also scales in a predictable and flexible way because you pay for what you need as you grow and avoid buying more than you need.

2. Ultrafast Streaming Performance

Panasas storage can support even the highest aggregate data rate requirements in your computing infrastructure and helps research organizations avoid the common and outdated practice of replicating data from shared storage to the Linux compute farm. That means accelerating time to publishing scientific results, a surefire path to success in the highly competitive world of life sciences.

3. Enterprise-Grade Reliability and Availability

In most NAS offerings, scalability comes at the cost of performance and reliability. With the ActiveStor solution, that trade-off no longer exists. Data reliability, performance, and availability actually increase as you scale Panasas storage, and you can use enterprise data services such as snapshots and quotas at no extra cost.

4. Management Simplicity for Lower Operating Costs

You can reduce storage complexity in the data center by consolidating many different storage products to as few as possible. Panasas offers the fully integrated ActiveStor scale-out NAS appliance with a single point of management, regardless of scale. As a result, one IT administrator can easily manage dozens to hundreds of terabytes of ActiveStor storage to start, and administration stays just as easy when the ActiveStor solution scales to petabytes of storage capacity.

What to Do Next

To learn how Panasas storage can drive innovation in your life sciences organization, visit www.panasas.com.

Notes

1. Stephen Davies, Michael Shanler, "Predicts 2016: Digital Generates Business Value Opportunities in Life Science," Gartner (December 2015), www.gartner.com/doc/3174119/predicts—digital-generates-business.
2. Chris Cestaro, "Predictions for Life Sciences Commercial Insights in 2016," Zephyr Health (December 2015), https://zephyrhealth.com/predictions-for-life-sciences-commercial-insights-in-2016/.
3. John Russell, "25 Percent of Life Scientists Will Require HPC in 2015," HPCwire.com (May 18, 2015), www.hpcwire.com/2015/05/18/25-of-life-scientists-will-require-hpc-in-2015.
4. "Challenging the Status Quo on Maintenance Contracts and Refresh Cycles to Lower Costs," Forrester Consulting (May 2013), http://global.networkhardware.com/rs/networkhardware/images/Challenging_the_Status_Quo_on_Maintenance_Contracts_and_Refresh_Cycles_to_Lower_Costs.pdf.