

Powering Modern Data Analytics

5 Ways to Help You Gain Better Insights Faster



Table of Contents

Introduction	3
5 Ways to Help You Gain Better Insights Faster	
1. Analysis of Real-Time Data	4
2. Simple, Agile Technology	5
3. Cloud Optimized	6
4. Modern Consumption Models	7
5. Continuous Application Uptime	8
Checklist	9
Pure Flashblade Is Made for Real-Time Analytics	10

Unleash the Power of Modern Data Analytics

A modern set of analytics applications is here, fueled by the global shift toward digital transformation. These applications take advantage of the massive amounts of data being generated every minute to deliver insights that improve operational efficiencies and user experiences; enhance compliance, regulatory, and security processes; and even drive new products and revenue streams.

Of course, it's more challenging than it sounds. Many organizations are oblivious to the vast amount of data, often spread across disparate sources, that is available to them. Or worse, they're unsure of how to tap into it all. Organizations need to analyze all of this data in real time, with a simplified process that's easy to manage.

In order to ensure that today's analytics tools deliver to their potential, it is critical that the underlying data platforms can keep pace. This guidebook outlines the five key criteria where your storage solution must deliver advanced capabilities to power today's modern analytics tools, from ingest to visualization, so that you can make better decisions faster.



“ **In order to ensure that today's analytics tools deliver to their potential, it is critical that the underlying data platforms can keep pace.** ”

1

Analysis of Real-Time Data

Faster time to insight is the difference between treading water and riding the wave of improved revenue growth and profitability. The better your organization is at utilizing analytics in real time, the faster it can investigate and seize potential business opportunities. To ensure real-time insights, make sure your storage solution has these characteristics:

High Throughput

Scalable, high-throughput systems ingest data faster for real-time analytics. As the amount of real-time data that is being analyzed increases, it is important to have a platform that scales throughput and capacity.

High Concurrency

Make sure you can support large numbers of users with massively concurrent connections from workstations, nodes, and database servers so that you can get data into the hands of more data scientists and decision-makers.

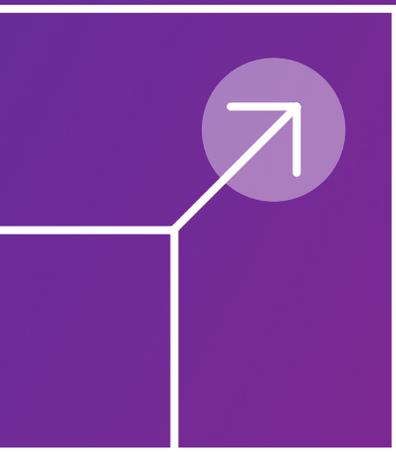
Consistent Performance at Scale

Look for linear, predictable storage performance while scaling up to tens of billions of files and objects. With high-performance all-flash systems like Pure FlashBlade®, you can reduce the time spent on searching, indexing, and analyzing data at any scale.

2

Simple, Agile Technology

Data architects face a constant flux of workload changes and operational requirements that shift from day to day and even hour to hour. Unless you look forward to constantly fine-tuning a system to specific workloads, make sure your storage is agile enough to handle whatever you want to throw its way—and more.



Support for Multiple, Variable Data Patterns

All-flash solutions like FlashBlade offer predictable, ultra-fast performance for any access pattern, whether random or sequential. From small, metadata-heavy files to large streaming files, your storage platform needs to be adaptable to modern architectures. FlashBlade is “tuned for everything,” offering performance for variable data without the need for manual tuning.


Ability to Scale With Multiple Workloads

A solution that disaggregates compute from storage will provide your organization with the most efficient platform for hosting multiple analytics applications, concurrently supporting a large number of users and easily scaling with data growth, allowing you to eliminate the complexity of siloed approaches. FlashBlade non-disruptively scales performance and capacity by simply adding blades.

3

Cloud Optimized

Simplicity and elasticity are what makes public cloud so attractive. You can enhance modern data analytics with an on-prem storage foundation that offers cloud benefits with on-premises ownership and control. Here's what to look for in cloud-optimized storage:



Ability to Adapt to Modern Cloud Architectures

Virtually all analytics software vendors have begun a shift to cloud-enabled object architectures such as Splunk SmartStore and Vertica in Eon Mode. Protect your infrastructure investments with a storage solution that is built to work for modern environments.

Consistent Architectures and Technology for Public and Private Clouds

Your storage solution should deliver cloud-like simplicity and agility with consistent performance and control for cloud-native and cloud-focused workloads in your own data center. Gain the ability to leverage common tools and skill sets across public and private clouds.

Fast Object Storage

Pure developed the industry's first native, unified, and fast file and object platform with FlashBlade. Internal tests demonstrate that FlashBlade speeds insights by delivering analytics teams radically faster query performance than traditional object storage platforms that disaggregate storage and compute.

4

Modern Consumption Models

Your choice of resource consumption should provide you with the best TCO. With a cloud-like consumption model, pay-as-you-go billing options provide you with the flexibility to scale up and down so you can consume IT resources to fit your business needs.

Subscription to Innovation

A subscription model like Pure Evergreen™ program simplifies purchasing and upgrades, keeps your platform modern, and allows you to free your budget for areas critical for analytics (such as compute).

True OPEX Model*

To keep operating expenses low, look for storage you can implement via a flexible consumption model that delivers easy installation, easy upgrades, and no performance tuning requirements. Pure-as-a-service offers a 100% OPEX service subscription model. With Pure, you have a choice of CAPEX and OPEX consumption to match the needs of your business model.

*OPEX treatment is subject to customer's auditor review.

5 Ways to Help You Gain Better Insights Faster

5

Continuous Application Uptime

Downtime—planned or otherwise—is bad for business. Look for storage designed to maintain application uptime.

Uninterrupted Performance

The right storage solution will maintain application uptime with non-disruptive performance through maintenance operations, software upgrades, and capacity expansion.

Disaggregated Scale-Out Architecture

This will give you the ability to add index nodes and rebalance data quickly, reducing node offline time during failures.

AI-Driven Support

AI-driven support will proactively identify and resolve issues before they become outages.

Easy Capacity Planning

Self-driving, automated storage management will simplify planning for capacity and performance needs.

Checklist

Does your storage solution have what it takes to power modern data analytics tools?

✓ **Analysis of Real-Time Data**

Power rapid insights from log analytics and more with scalable, all-flash storage.

✓ **Simple, Agile Technology**

Adapt to any workload with “tuned for everything” systems.

✓ **Cloud Optimized**

Get the flexibility and scalability of public cloud from your on-prem storage.

✓ **Modern Consumption Models**

Lower your TCO and free your budget for other areas critical to analytics.

✓ **Continuous Application Uptime**

The goal is non-disruptive performance through maintenance, expansion, and upgrades.



Pure Flashblade Is Made for Real-Time Analytics.

One storage solution covers all the criteria necessary for meeting the demands of today's modern data analytics applications. You can improve the quality and quantity of your insights by running your analytics workloads on fast file and object storage like FlashBlade from Pure Storage.

Accelerate Insights

FlashBlade gives you the advantage of linear, predictable storage performance while scaling up to tens of billions of files and objects. You'll achieve massive throughput of up to 15GB/s per chassis and faster search and indexing of data, all while supporting massively concurrent connections to/from workstations, nodes, and database servers.

Simply More Agile

FlashBlade is "tuned for everything," so you can easily adapt to workload changes, from batch to real time, sequential to random I/O, or small to large file size. It lets you support multiple, disparate workloads with non-disruptive multidimensional scale-out.

Ahead in the Cloud

Combine the control and efficiency of on-premises solutions with cloud-like economics. FlashBlade fast-object storage allows you to run cloud applications with consistency and performance not possible from alternative solutions.

Evergreen Is Pure Gold

The Pure Evergreen subscription model simplifies purchasing and includes automatic upgrades, allowing you to focus resources on other areas critical for compute. Evergreen is a subscription to innovation that provides superior choice for resource consumption and the only storage TCO that goes down over time.

Always on and Available

FlashBlade is designed to maintain application uptime. You'll get 100% non-disruptive performance through maintenance operations, software upgrades, and capacity expansion. Node offline time is reduced by 94% during failures, and AI-driven support allows Pure's support teams to proactively identify and resolve issues before they become outages.

Proven and Recognized

A customer experience NPS score of 85.2 puts Pure in the top 1% of all businesses. That's just one of the reasons the Gartner Magic Quadrant has named Pure the storage leader five years in a row.

Storage That Makes You Smarter

Pure FlashBlade powers better, faster decisions. If your organization is ready to activate modern analytics, start with the architecture built to meet the demands of today's workloads.

[Click here to begin](#)

