WHITE PAPER

Power Modern Applications with a New Approach to Data

Today's enterprises are turning to unified fast file and object (UFFO) storage to maximize the value of unstructured data.

Contents

Enterprises run on data, and the biggest wins go to organizations that know how to collect, store, manage, and analyze that data most effectively. Making this happen requires harnessing critical insights from an ever-growing number of data sources, both structured and unstructured.

In the beginning, computer data was structured and fairly straightforward to store and manage. Information fit into relational databases, like Excel spreadsheets and SQL databases, with predefined data models. Some of this data remains structured, which is easy to search and manage and lives on the hard drive as structured files.

Back in the day, we also had unstructured data in the form of content repositories, file-shares, and home directories. However, as computing technology continues to evolve, new unstructured data workloads are emerging that don't fit into the simple table structure. This unstructured data includes audio files, images, video footage, sensor data, and social media posts. It also includes text files and documents that need context to parse—think notes typed into a customer service application during a call or a transcript of witness testimony in a court case.

The challenge of all these new data types is multifaceted. Efficiently storing and simplifying the management of this process is mandatory. However, searching and analyzing unstructured data—and extracting value from it—is much more complex than working with structured data. And the worldwide volume of unstructured data is exploding. In fact, industry experts project that by 2025, 80% of all data created will be unstructured.

While data has the ability to generate a bounty of meaningful insights, informed decisions, and innovations, the proliferation of new data types and sources, data growth, and manageability must be governed effectively. In other words, if enterprises want to thrive in today's fiercely competitive marketplace, they require a more efficient and effective way to store, manage and leverage unstructured data.

Today's Applications Can Create Breakthroughs—for Some

Unstructured data volumes are growing exponentially as enterprises increasingly deploy advanced applications to gain breakthroughs in product development, increase revenue models, and deliver more personalized customer service. According to Bredin research, 63% of enterprise organizations deployed or increased the use of customer-facing applications in 2020.[1] The rise of cloud storage and technologies such as containers has coincided with the growth of cloud-native applications and changed the way applications today are developed and deployed. However, achieving measurable value from these applications requires more than just deployment. The right infrastructure is necessary to maximize the potential of the application, analyze the data efficiently, and discover innovative insights from it. Legacy data solutions and architectures are simply not up to the challenge. Applications that tap into the internet of things (IoT) or employ artificial intelligence (AI) and machine learning (ML) all create massive volumes of unstructured data. So do applications that assist in activities such as fraud detection, predictive fulfillment, instant search, or autonomous driving. These are just some examples. In fact, all of us interact with dozens of applications on a daily basis which is driving the unstructured data growth.

Unstructured data needs to be processed and analyzed before you can gain those valuable insights. And processing and analyzing it creates even more data.
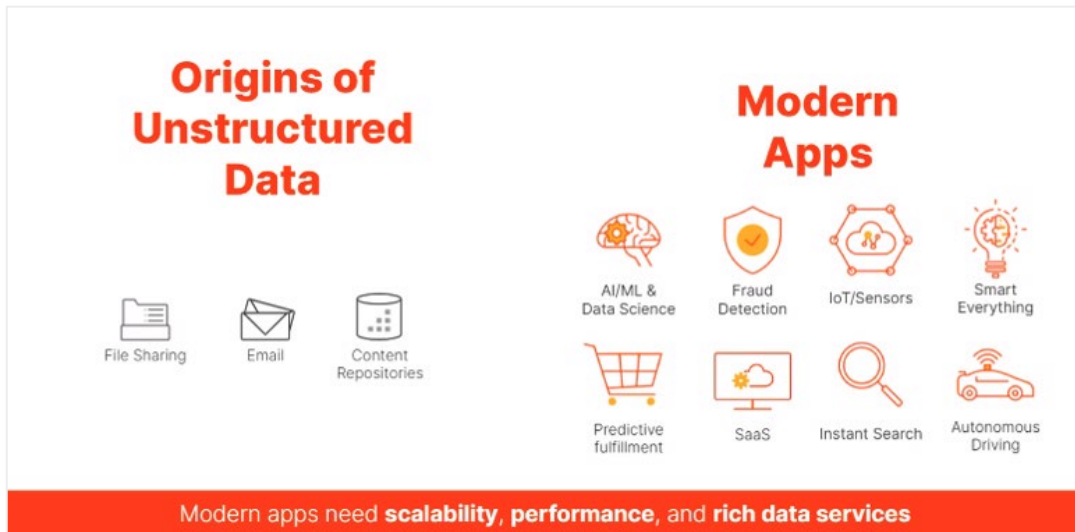
The challenges in storing, managing, and analyzing unstructured data come from—well, its lack of structure. Traditional infrastructure approaches often require your data from disparate applications and systems to be stored in dedicated silos across your infrastructure. The information isn't managed or accessed in any unified way. Unstructured data can be stored as files or objects in a variety of different storage environments, but this makes it difficult if not impossible to find or access when you need it.

Getting maximum value out of your unstructured data is all about getting deep into its many layers. Fast, efficient analysis is the key. That means consolidating data from multiple sources and locations. Legacy storage systems just can't deliver the flexibility, speed of access, or unified management you need to adequately handle unstructured data and make it deliver. That requires a new approach and a different kind of infrastructure.

That's a unified fast, file and object storage platform.

UFFO Storage: A New Approach to Meet Evolving Needs

Unified fast file and object (UFFO) storage is a new category of storage that addresses the needs of modern data types and applications. It is a single, unified storage platform optimized for data that spans capacity-optimized and performance-optimized file and object workloads. UFFO storage integrates file storage (which delivers fast access) with the supreme scalability of object storage, an architecture that stores data as self-contained units called objects. It allows you to manage and store both types of unstructured data in a single, high-performance scale-out solution. With UFFO storage, you can eliminate siloed data warehouses, as well as legacy object stores, network-attached storage (NAS), purpose-built backup appliance (PBBA) and direct-attached storage (DAS).

It provides the right balance between, performance, simplicity, and agility for unstructured data management and storage for today's most demanding applications, including large-scale analytics and AI/ML workloads. UFFO storage can match the processing power of the most advanced chipsets to deliver maximum throughput.

With UFFO storage, you get multidimensional performance with high throughput, regardless of workload, file size, protocol, or file and object counts. The unified platform easily scales to support multiple applications, unlike single-use silos that separate file and object data. The solution also enables you to fine-tune utilization by decoupling compute and storage and allowing them to be scaled independently on demand.

UFFO storage can be the smart foundation of your modern applications, which need to be:

- Extremely flexible and dynamic and optimized to run across a large number of containers
- Distributed and easily scalable across your entire infrastructure
- Automated and orchestrated through APIs
- Resilient to always-evolving security threats and downtime due to planned or unplanned outages
- Developed to process and use data in real time

The most important takeaway here is that UFFO storage enables your advanced applications and analytics solutions. It transforms unstructured data into more than an unexamined by-product of an IoT device or a load of information that's simply archived for compliance reasons. UFFO storage empowers you to derive real-time insights that can truly transform your business.

Those Who Unlock Today's Data Can Become Tomorrow's Leaders

The rewards are great for those who can extract maximum value from their data. As more enterprises find smarter ways to use and learn from their data, customers will increasingly come to expect rich, immersive experiences that keep them coming back.

Embracing the modern approach to data and its insights can position your organization as a leader among competitors. Benefits of a UFFO storage platform include:

Simplicity
Take the complexity out of data storage to improve IT productivity. Gain cloud-like agility with simple, non-disruptive data management.

Performance
Deliver outstanding speed and performance across all of your file and object workloads to enable today's (and tomorrow's) modern applications.

Consolidation
Eliminate inefficient, sprawling silos of legacy applications and consolidate data for modern applications on a single platform to improve ROI and protect your investment.

How Pure Storage Can Help

At Pure Storage®, we believe in the power and efficacy of a UFFO storage platform. In fact, we're proud to offer the industry's most advanced all-flash storage solution for consolidating fast file and object data: FlashBlade®.

Architected to address the demands of modern applications and modern data, FlashBlade delivers the simplicity and multidimensional performance that enables the consolidation of key unstructured data workloads on a single platform.

Traditionally, unstructured data storage systems have been either highly scalable for storage but very slow, or fast but only meant for smaller batches of data. And both options handled only one type of data, file, or object.

FlashBlade, the leading UFFO storage platform, is something different. Its architecture is massively parallel and all-flash, for the highest performance. And it manages both file and object data without a hassle. Scaling out for improved capacity and performance is a matter of simply adding more blades.

Andrew McArthur, Ph.D., is in charge of the McArthur lab at McMaster University. He chose FlashBlade to accelerate decision-making in the fight against life-threatening superbugs and viruses. Using the UFFO platform to handle highly complex processes associated with massive genomic data sets, the lab has sped up the time to research and seek cures. And as COVID-19 hit, it was able to pivot immediately and leverage the high-performance infrastructure to help combat the virus.

"There's no point in playing with traditional storage because it's just not fast enough," says McArthur. "With Pure Storage, we can stay ahead of the curve as we fight global threats to human health."

With FlashBlade, you can also get powerful support for heavy-duty data analytics and data protection use cases, such as rapid restore, ransomware mitigation and remediation. You also get support for a wide variety of other use cases, including AI/ML, cloud-native applications and architectures, and more.

The platform's unique design includes a unique flash translation layer that allows the system to connect to each flash chip, creating a massively parallel data path to accelerate concurrent data access for today's data-intensive applications. The platform eliminates the use of serial protocols associated with using disk software to talk to flash. The solution is designed for immense but simple scalability. And the solution is powered by an innovative converged fabric that eliminates the complexities associated with networking when using traditional storage solutions. All of this allows you to modernize your applications, meet the always-evolving demands of today's modern applications, and eliminate bottlenecks for faster time to market and greater return on investment.

Learn more about UFFO and see how FlashBlade can help.