

ParallelStor

INTRODUCTION

It is a PFS storage solution specially developed for massive enterprise requirements with high availability as one of its many highlights. It is a proven integrated solution with a simple and scalable infrastructure to keep abreast with the most demanding workloads. The solution boasts of some brilliant performance metrics such as high throughput, high IOPS, high bandwidth (EDR-100GB/s, HDR-200GB/s), and high density (in various configurations).

WHY USE PARALLELSTOR?

With the proliferation of new and emerging technologies such as High-Performance Computing (HPC) artificial intelligence (AI) and containers the requirement for specialized storage to take care of specific performance has been on the rise. The requirement from new applications is primarily concerning high IOPS (input/output operations per second), and parallel accessibility of data without degrading any node performance. The need of the hour is a high-performance PFS storage solution like ParallelStor.

The main benefits are True Parallel Access, Highly Scalable, Low TCO, Enterprise Ready with Fault Tolerance, High Availability, Simple Deployment & Easy Management, and Minimum Metadata Access Latency.

MAINTAIN CUTTING-EDGE CAPABILITIES FOR EXTREME WORKLOADS

Storage is as critical as incredible computing power. For example, even the fastest and leanest machine will not draw applause if the storage is not equally fast.

ParallelStor understands the conundrum and provides consistent concurrent, real-time access to your data. To prevent bottlenecks and facilitate continuous high-performance data spread across multiple servers and their backend storage. It also has an in-built technology to monitor workloads and automatically adjusts configurations for maximum performance.



MAIN BENEFITS



TRUE PARALLEL ACCESS

Through its inherent design, ParallelStor separates file content and metadata. While storage servers take care of storing portions of the actual file content, the metadata server helps in the coordination of file placement and striping. The client can directly connect with storage servers and communicate with more than one server, leading to true parallel access to the file data.



LOW TCO

Apart from the deployment cost, one must also deliberate the operating costs of the solution once deployed. ParallelStor HDDs and SSDs adheres to strict performance, quality and interoperability requirements. The additional scrutiny considerably reduces the failure rate eventually bringing down the total cost of ownership.



MINIMUM METADATA ACCESS LATENCY

ParallelStor is also capable of distributing the metadata across multiple servers to minimize the metadata access latency. By doing so, each of the metadata servers stores only a section of the global file system.



ENTERPRISE READY WITH FAULT TOLERANCE AND HIGH AVAILABILITY

ParallelStor delivers a high-availability (HA), shared disk solution for PFS storage. The solution introduces an additional level of data access protection without the requirement for data duplication. The protection ensures an enterprise ready PFS which instantly recovers from any unseen fault or failure, keeping the crucial application always running.



PROMPT RESPONSE

To meet the demands of high-volume data applications such as AI, and 3D imaging, you need a storage solution that instantly respond to data growth. The scalability and flexibility of ParallelStor enables to adjust to the evolving workloads by responding quickly.



SIMPLE DEPLOYMENT AND EASY MANAGEMENT

In an environment where complexity is rampant, simplicity is critical. Unlike the existing parallel file system in the market, ParallelStor does not need multiple platforms and models to scale. The storage solution is easy to deploy and support.



SCALABLE

From its very inception, ParallelStor has been optimized for a high data throughput with a concerted effort on flexible scaling. Technically, the solution encompasses multiple storage servers to deliver a highly scalable and flexible network file system that supports striped file content.



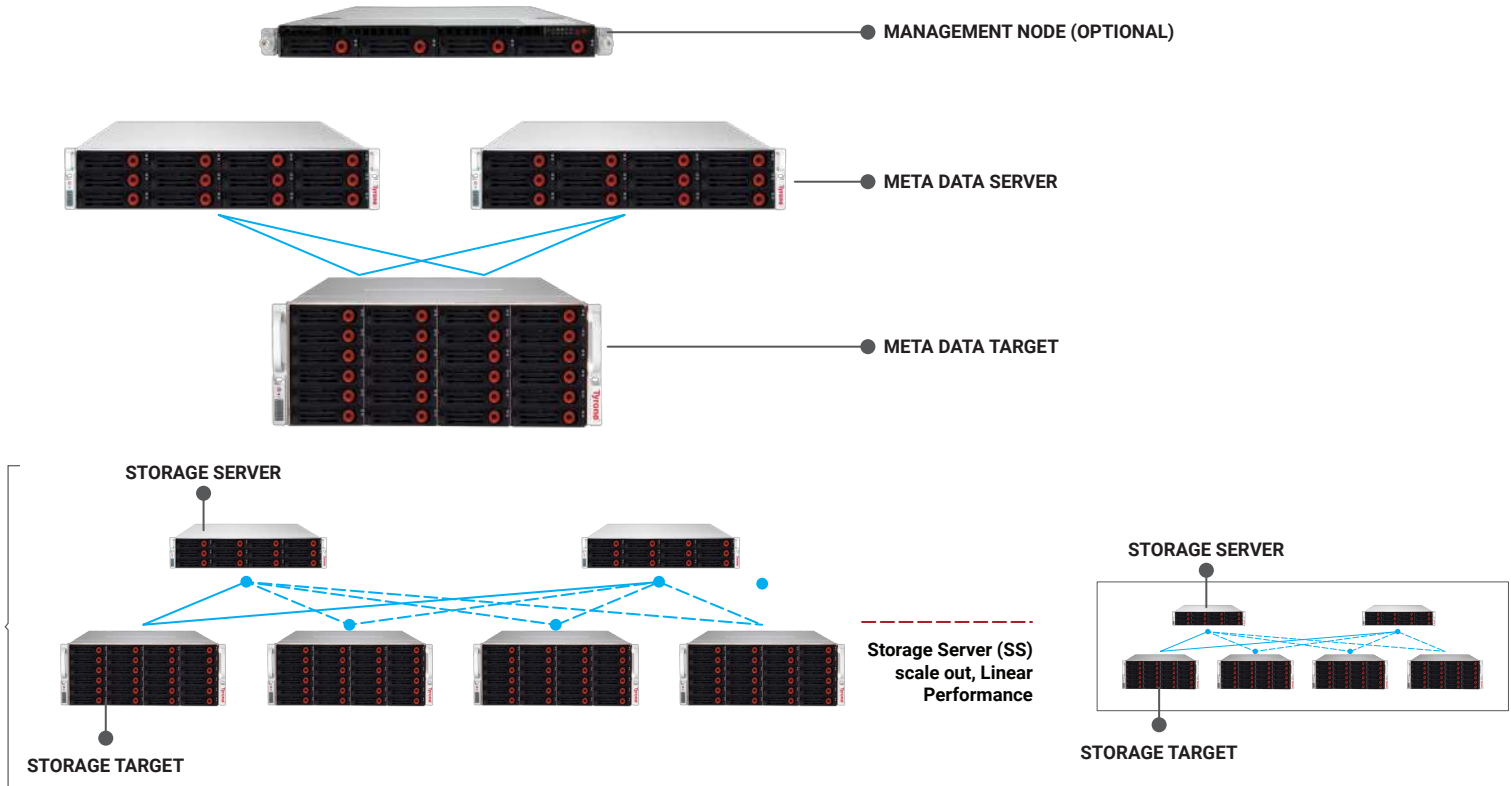
KEY FEATURES OF PARALLELSTOR

- Supports large datasets and high input/output operations per second (IOPS) requirements
- Large, fast distributed scratch file system
- A centralized storage for clusters
- Simple building block architecture delivers predictable scaling to specific requirements
- Industry-leading storage density with High Availability design
- Total Solution including Systems, Software, and Services
- InfiniBand, or Ethernet as a high-speed interconnect with dynamic failover capability between different network topologies
- No single point of failure
- Reduce storage costs by up to 90%
- Flexible data striping per directory or file

WHERE IT IS BEING USED



PFS STORAGE SOLUTION FOR HPC WITH HIGH AVAILABILITY (HA)



**Note: The diagram is for representation purposes only and the actual component might differ in a real-world scenario.*

Tyrone Bespoke PFS Storage Solution for HPC or ParallelStor is a wholly supported storage solution with high throughput. The solution includes a management node - that is optional - a pair of metadata servers (MDS), a metadata target (MDT), a pair of storage servers, and related storage arrays. The pair of MDS is connected to MDT by high-performance links. The pair of storage servers host the storage target for Parallel File System.

The PFS Storage Solution Architecture Consists of Four Main Services:



Management Service: ParallelStor management service is set up on both metadata servers. The management store is initialized directly in the management directory on the metadata target.



Storage Service: In a high-capacity, ParallelStor solution the data storage is achieved across four storage targets or arrays. On each storage array, linear disk groups of variable drives are created. For every disk group, a single volume utilizing all the space is also created.



Metadata Service: Within ParallelStor, every metadata service takes care of only one metadata target. For example, if there are twelve MDTs, there need to be twelve instances of metadata services. The metadata targets are formatted to perform exceptionally with small file operations. Moreover, ParallelStor also captures information directly on the inodes of the file system for delivering optimal performance.



Client Service: The ParallelStor client module is loaded on all the hosts that need access to the file system. Once the module is loaded and the client service is started, the service mounts the file system predefined in the module. With the approach, the client service starts similarly to other Linux services via the service start-up script and facilitates the automatic recompilation of ParallelStor client module once the system updates.

Tyrone[®]

ParallelStor

CONTACT US

E-mail info@tyronesystems.com

Website www.tyronesystems.com



facebook.com/tyronesystems

twitter.com/tyronesystems

linkedin.com/company/tyrone-systems

www.tyronesystems.com

