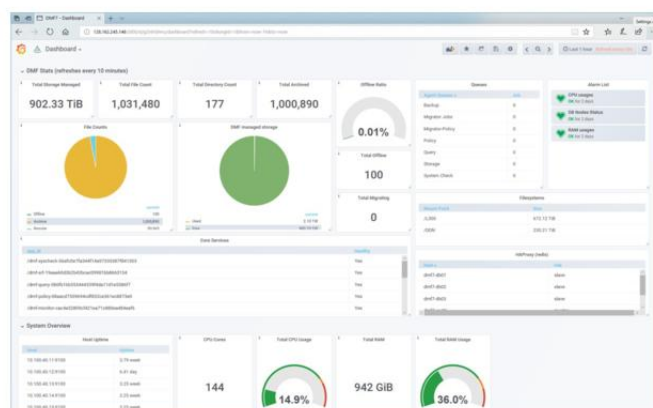


# HPE DATA MANAGEMENT FRAMEWORK 7

## High Performance Computing Storage Solutions

### HPE Data Management Framework Monitoring GUI



## WHAT'S NEW

- A scalable database. The data base is distributed across nodes for higher system availability, is scalable to billions of files, and delivers high performance for queries.
- Namespace reflection. Maintain a reflection of managed file systems and use it for data management operations instead of compromising the performance of the file system metadata servers.
- Support for extensible metadata.

## OVERVIEW

Are your High Performance Compute (HPC) and AI environments struggling with file management? HPE Data Management Framework 7 (DMF7) delivers centralized data management across HPC and AI storage systems and protects scalable, parallel file systems like Lustre and Spectrum Scale. Namespace reflection is used to create an independent snapshot of file system state, allowing you to recover file systems in a known good state. This system maintains file versions, allowing users to recover files from previous successful job runs.

---

Administrators can use metadata to customize queries, use extensible metadata for policy management and create personalized data sets from associated files.

- Data curation. Users can stage their files into an independent back end data base for long term curation. Later, users can locate and recall their files to an entirely different file system.
- Query engine. Administrators, users, and the built-in policy management engine all use a common query tool to locate files and filter the results.
- Native file system integration. Native integration uses the standard tool set for Lustre and Spectrum Scale to increase utilization and automate data management operations.

HPE DMF7 automates data movement between tiers in a storage hierarchy, e.g. between flash and disk. Administrators and users can also use HPE DMF7 to move files between file systems, e.g. when files must be moved from storage that is being retired. HPE DMF7 improves utilization of expensive, high performance storage by automatically moving files to lower cost storage tiers, creating a virtual storage space that appears to scale beyond the physical capacity.

## FEATURES

### **HPC/AI Applications are Data Rich, and Data Must be Protected**

HPE Data Management Framework 7 creates immutable versions of files and takes snapshots of the namespace reflection. Managed via an administrator policy, recovery can be customized for lowest RTO from disk, lowest cost from tape, and/or from remote locations via S3/cloud.

Loss of the file system due to failure has catastrophic impact upon availability of the high performance compute cluster. Even when a file system has tools for repairing, the complexity and time it takes to repair a broken file system can extend the compute outage beyond acceptable SLAs.

Up to now, protecting file systems and data has been a costly investment with imposing drawbacks, including the lack of backup windows, backup utilities that are sub-optimized, optimized for PB-sized parallel file systems, and the negative impact upon performance from scanning file system metadata.

HPE DMF7 creates immutable versions of files and takes snapshots of the file system state via namespace reflection. Both operations are managed via administrator policy. Recovery can be optimized for lowest RTO from disk, lowest cost from tape, and/or from remote locations via S3/cloud.

### **The Emergence of Exascale Computing is Challenging the Scaling Limits of Legacy HPC Storage**

HPE Data Management Framework 7 manages free space in storage by automatically moving 'stale' files out of the high performance storage, creating an underlying storage space that is bigger on the inside. Administrators easily manage policy settings, alleviating the need to take brute force actions.

Perhaps the challenge is felt most acutely by storage administrators, who struggle to maintain enough free space in costly high performance storage while users independently flood the file system with new files. Increasing the storage budget and/or deleting user files aren't practical remedies.

The volume and diversity of data demanded by HPC/AI applications has fueled the growth of the "storage beast" that feeds on HPC budgets. At the same time traditional parallel file system architectures are struggling under the weight of relentless growth in the number of files and inodes.

Eventually, administrators need to work with users to prune unused files from the file system to ensure metadata performance isn't undermined. When old files are marked for removal, no data has to be moved since HPE DMF7 already preserves files and metadata in less costly back end storage.



### **HPC/AI Storage Environments are Diverse and Data has to be Portable**

HPE Data Management Framework 7 automatically migrates files down the storage system hierarchy without administrator interaction and recalls them up to high performance storage on demand. It uses parallel data movers and the high speed network to move files faster than standard desktop utilities.

Managing HPC/AI data movement is an intimidating task. Tools aren't easy to use, they don't scale well, network pipe bandwidth is limited, and users may not have the needed skills. When data cannot be moved easily and the motivation to move it is low, the default choice is to leave it in place.

Storage systems are optimized for performance, capacity, and cost, and data is always in flight between these tiers. Application workflows demand that data follow the user and the application, and administrators are continually pressured to manage storage costs and push data down the hierarchy.

Technology migration is a common driver of data movement and HPE DMF7 future proofs against this risk. It automates the migration of back end objects from older, inefficient generation HDD/tape technologies and on to generations that have the highest density, reliability and performance.

When the file system must be retired, administrators move the file system and files onto HPE DMF7 back end devices. Once they are protected, the file system can be staged into an entirely new namespace and files can be staged into the new namespace or remain in curation by HPE DMF7.

### **Reducing HPC Storage Costs Means More Budget for Compute**

HPE Data Management Framework 7 is the HPC storage management platform that automates data workflows and reduces HPC storage costs, so HPC customers can spend their valuable project budgets on the infrastructure resource that matters most.

The storage beast works against that goal. The storage beast encourages HPC customers to replicate files on expensive proprietary storage, scale out the most expensive storage tier to satisfy data growth, and says it's easier to leave data in place on expensive storage than it is to move data.

The primary goal of HPE DMF7 is to reduce storage costs. Instead of purchasing more expensive storage, HPE DMF7 makes it easier for HPC customers to protect, scale and move data using lower cost storage and improve utilization and performance of mission critical HPC storage resources.

The unabated growth of HPC data and the orchestration of large data sets for AI/machine learning are driving unprecedented growth in storage capacity needs. And the adoption of flash storage at a higher effective cost means that containing cost is still the primary goal for HPC storage buyers.



For additional technical information, available models and options, please reference the [QuickSpecs](#)

## HPE POINTNEXT SERVICES

[HPE Pointnext Services](#) brings together technology and expertise to help you drive your business forward and prepare for whatever is next.

### Operational Services from HPE Pointnext Services

[HPE Pointnext Tech Care](#) provides fast access to product-specific experts, an AI-driven digital experience, and general technical guidance to help enable constant innovation. We have reimagined IT support from the ground up to deliver faster answers and greater value. By continuously searching for better ways to do things—as opposed to just fixing things that break—HPE Pointnext Tech Care helps you focus on achieving your business goals.

[HPE Pointnext Complete Care](#) is a modular, edge-to-cloud IT environment service that provides a holistic approach to optimizing your entire IT environment, and achieving agreed upon IT outcomes and business goals through a personalized and customer-centric experience. All delivered by an assigned team of HPE Pointnext Services experts.

**HPE Integration and Performance Services** help you customize your experience at any stage of your product lifecycle with a menu of services based on individual needs, workloads, and technologies.

- Advise, design, and transform
- Deploy
- Integrate and migrate
- Operate and improve
- Financial Services
- Greenlake Management Services
- Retire and sanitize
- IT Training and personal development

### Other related services

[HPE Education Services](#) delivers a comprehensive range of services to support your people as they expand their skills required for a digital transformation. Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and support options.

**Defective Media Retention** is optional and allows you to retain Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.

## HPE GREENLAKE

[HPE Greenlake](#) is HPE's market-leading IT as-a-Service offering that brings the cloud experience to apps and data everywhere – data centers, multi-clouds, and edges – with one unified operating model. HPE GreenLake delivers public cloud services and infrastructure for workloads on premises, fully managed in a pay per use model.

If you are looking for more services, like **IT financing solutions**, please [explore them here](#).

Make the right purchase decision.  
Contact our presales specialists.

[Find a partner](#)



Chat now (sales)



Call now



Buy now



Share now



Get updates



Hewlett Packard  
Enterprise

© Copyright 2022 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Parts and Materials: HPE will provide HPE-supported replacement parts and materials required to maintain the covered hardware.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

Linux® is a registered trademark of Linus Torvalds.

Image may differ from the actual product  
[PSN1010144088TWEN](#), May, 2022.