



The Blue Whale File System (BWFS) is a media-centric distributed file system developed and customized for media applications. Through virtualization, it simplifies the management of storage resources, offering tremendous performance, scalability, and concurrent high-speed file access to heterogeneous clients, all across a single global namespace. BWFS enables organizations to efficiently create, manage, and deliver content throughout their digital workflows, addressing the different storage challenges in today's media-rich era.

# Blue Whale File System

## Media-centric distributed file system

### Key benefits

#### Scalability

- Massive scalability (up to 64ZB for massive content storage and rapid expansion capability)
- Non-disruptive scaling for better productivity (no need to stop or delay production or playback when adding more storage)

#### Performance

- Support more simultaneous SD or HD streams for playback or editing sessions
- Simultaneously support different applications (i.e. playback, editing, etc.) through efficient bandwidth management
- Better QoS through client bandwidth control and file system tunability

#### Resource Management

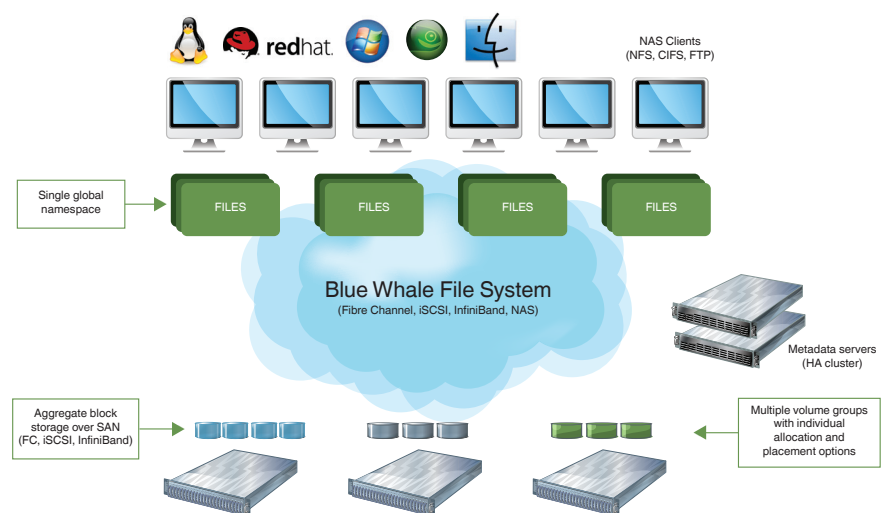
- Optimized capacity utilization (i.e. free space across the entire file system is consolidated and made easily available, no underutilized storage systems)
- More efficient workflows by reducing file transfers between parts of the workflow or applications with different operating systems

#### Cost

- Better ROI and reduced acquisition costs through efficient capacity utilization
- Lower operating costs through simplified management
- Hardware independence means more vendor choices and flexibility in cost
- Pay as you grow

The Blue Whale File System consolidates different storage systems into a single storage pool where all connected clients share a single global namespace. It supports a massive capacity of up to 64ZB, billions of files, and hundreds of clients, all in a single file system. Multiple volume groups can be defined with individual data allocation and placement parameters, allowing ultimate flexibility in supporting different requirements of connected applications. Storage reconfiguration is also transparent to clients, requiring no client configuration changes or downtime.

BWFS features a high-performance, low-latency architecture that provides direct block-level access simultaneously to BWFS clients. Employing an out-of-band data transfer architecture, data is directly transmitted between connected clients and physical storage over Fibre Channel, iSCSI, or Infiniband, allowing maximum performance and overcoming the long latencies and overhead of traditional file systems.



### Key Features

#### High real-time performance

BWFS employs configurable write behind and read ahead buffers to provide high real-time performance.

#### Media-centric file layout

With a media-centric file layout, BWFS allows space efficiency for small files and high streaming performance for large media files.



**File change notification**

Unlike other file systems, BWFS features file change notification that allows media applications to maintain up-to-date content inventory without re-scanning massive number of files.

**Client-based bandwidth control**

BWFS provides secure bandwidth control that is more aligned with how customers want to manage storage access in a multi-user environment vs volume-based bandwidth control offered by other vendors. Bandwidth caps can be set on each client through a GUI utility in order to ensure QoS to high-priority real-time applications and limit bandwidth usage of non-priority users.

**Highly tunable**

BWFS provides many tunable parameters and configuration knobs to enable tuning and customizing for media applications and underlying storage systems. This ensures that performance is optimized according to application behavior and requirements.

**Dynamic non-disruptive scaling**

Storage expansion is nearly instant, allowing a storage grid to scale either vertically or horizontally in less than 10 seconds. Storage expansion also happens online, without disruption to client access, and without impact to client performance.

**Access control**

BWFS is fully integrated with LDAP and Active Directory for IT-friendly multi-user access administration.

<b>Maximum Capacity</b>	64ZB per system
<b>Maximum number of file systems</b>	16 per MDS
<b>Maximum file size</b>	32TB
<b>Maximum number of files</b>	Billions of files depending on file size
<b>Maximum number of LUNs</b>	4096
<b>Maximum LUN size</b>	Variable – limited only by maximum drive logical block addressing
<b>SAN technology support</b>	Fibre Channel, iSCSI, Infiniband
<b>NAS protocol support</b>	CIFS, NFS, FTP, HTTP
<b>Client OS Support</b>	<ul style="list-style-type: none"> <li>- Microsoft Windows XP/2003/2008/Vista/7 (32-bit &amp; 64-bit)</li> <li>- Red Hat Enterprise Linux 4.x/5.x (32-bit &amp; 64-bit)</li> <li>- CentOS Linux 4.x/5.x (32-bit &amp; 64-bit)</li> <li>- Oracle Enterprise Linux 4.x/5.x (32-bit &amp; 64-bit)</li> <li>- SuSE (32-bit &amp; 64-bit)</li> <li>- Mac OS X 10.5.x – Leopard (32-bit)</li> <li>- Mac OS X 10.6.x – Snow Leopard (32-bit &amp; 64-bit)</li> </ul>
<b>HA features</b>	<ul style="list-style-type: none"> <li>- Active/Standby Metadata Servers</li> <li>- Metadata journaling</li> </ul>
<b>MDS OS requirements</b>	<ul style="list-style-type: none"> <li>- Red Hat Enterprise Linux 4.x/5.x (32-bit &amp; 64-bit)</li> <li>- CentOS Linux 4.x/5.x (32-bit &amp; 64-bit)</li> <li>- Oracle Enterprise Linux 4.x/5.x (32-bit &amp; 64-bit)</li> </ul>
<b>MDS hardware recommendations</b>	<ul style="list-style-type: none"> <li>- Memory: 2GB RAM minimum (8GB recommended)</li> <li>- CPU: Dual-core AMD Opteron and Intel Xeon EM64T or higher</li> <li>- System disk: 80GB minimum disk space</li> <li>- Network Interface Card: Gigabit Ethernet network cards that are supported by Linux FC HBAs: QLogic 23xx, 24xx, 256x</li> </ul>

