

# EMC Centera Universal Access

## Expanded application support means simplified deployment and integration

### The Big Picture

- Purpose-built, networked software appliance for EMC Centera that greatly expands application support
- Fully compliant with industry-standard CIFS, NFS, FTP, and HTTP protocols to assure interoperability with the widest range of solutions
- High-performance and highly scalable caching file system: sustains efficient write performance, automatically synchronizes with EMC Centera, and supports up to 200 million files
- Retention period support: assures content availability and authenticity by implementing retention policies
- Enhanced availability: provides support for increasing availability of environments
- EMC developed, sold, and supported: provides a single point of contact for installation and support of EMC Centera Universal Access software

EMC® Centera® Universal Access is a networked software appliance that enables enterprise applications to work with EMC Centera® using industry-standard file system protocols such as NFS for UNIX, Linux and IBM System i (formerly iSeries and AS/400 applications) and CIFS (for Windows applications), plus Internet protocols such as FTP and HTTP. It is the only software appliance based on industry-standard protocols directly sold and supported by EMC as a purpose-built extension to EMC Centera. EMC Centera Universal Access is used when fixed-content applications are not integrated with the EMC Centera application programming interface (API).

EMC Centera Universal Access makes it possible for more enterprises to capture the industry-leading benefits of the EMC Centera storage platform immediately without modifying existing application environments. Based on a high-performance caching file system, EMC Centera Universal Access stores and moves data to or from EMC Centera and guarantees automatic synchronization with EMC Centera. In addition, EMC Centera Universal Access supports retention policies and offers an enhanced availability mode for increasing availability of environments.

With EMC Centera Universal Access, any enterprise application that can mount a network drive or use FTP and HTTP can immediately capture EMC Centera's unique benefits, such as long-term retention and assured, long-term content integrity. From home-grown applications to non-integrated versions of applications from EMC partners, EMC Centera Universal Access makes it possible to utilize EMC Centera in customer environments with no change to the existing application, greatly simplifying and accelerating deployment. As legacy applications are upgraded to natively integrated versions, EMC Centera Universal Access remains useful as a general-purpose access point for fixed content from applications that are not yet integrated with the EMC Centera API.

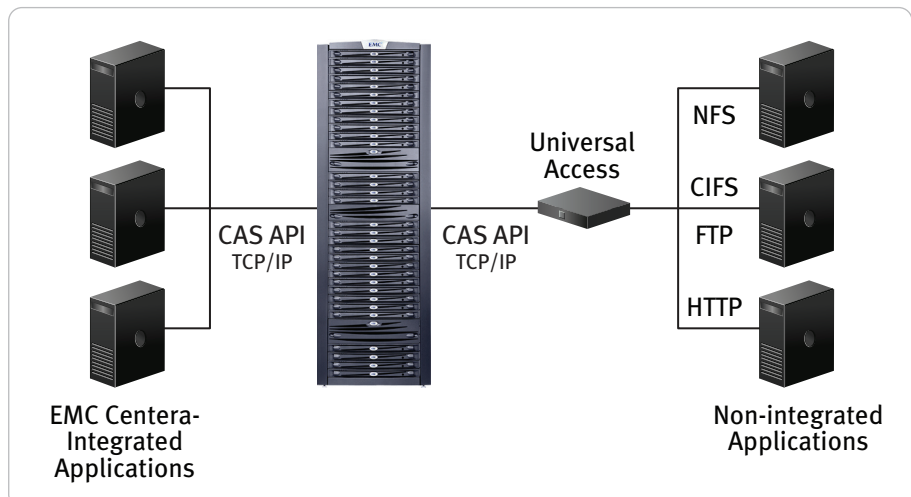
### A purpose-built extension of EMC Centera

Optimized for fixed content, EMC Centera Universal Access is purpose-built to handle large objects and read/write-intensive applications such as medical image archiving and document management. And with the ability to handle over 200 million objects in a single file system, preserve and honor data retention policies, and support enhanced availability configurations, EMC Centera Universal Access is an integral and trusted part of EMC Centera deployments in enterprises the world over. That's why EMC Centera Universal Access is the only software appliance based on industry-standard protocols developed, sold, and supported by EMC.

### Operating EMC Centera Universal Access

**EMC Centera Universal Access is a networked storage appliance that is able to:**

- Interface EMC Centera with any application over standard file interfaces such as NFS, CIFS, FTP, and HTTP from open systems and IBM System i clients
- Present to these applications the complete object store in EMC Centera
- Store/forward (cache) a subset of the data for EMC Centera



The basic functions of the EMC Centera Universal Access are:

#### Application-level Operations

- Store and modify files
- Retrieve files
- Delete files
- Retain files

#### Background Operations

- Synchronize with EMC Centera
- Flush file data and retain file metadata
- Back up and restore metadata
- Manually fail-over to a standby Universal Access and recover content

#### Store and modify files

Applications perform standard network file write operations to EMC Centera Universal Access via NFS, CIFS, or FTP. EMC Centera Universal Access must be in “the path” for all file creations and is responsible for maintaining the metadata and content addresses for all files. EMC Centera Universal Access is also responsible for holding the data portion of a file until the data is written to EMC Centera. Applications may also modify existing files stored in the EMC Centera Universal Access file system; EMC Centera Universal Access will forward the modified file to EMC Centera. EMC Centera Universal Access supports multiple access profiles and virtual pools enabling customers to choose management and replication policies on a per-application basis.

#### Retrieve files

EMC Centera Universal Access provides application access to all files that it stores on EMC Centera and maintains a subset of those objects in its local cache at any given time. It presents a complete file list to applications whether the object is actually stored on EMC Centera Universal Access’s local disk cache or only on the EMC Centera cluster. When an application requests a file via NFS, CIFS, FTP, or HTTP, EMC Centera Universal Access first checks to see if the file is stored in its local cache and if so, delivers the file to the application. If an application makes a request for a file that is not stored on EMC Centera Universal Access’s local cache, EMC Centera Universal Access fetches it from EMC Centera, delivers the file to the application, and also stores the file in the cache for future access.

#### Delete files

Applications perform standard network file delete operations on EMC Centera Universal Access via NFS, CIFS, FTP, or HTTP. Unless the data in question has been assigned a retention period, Universal Access will delete the file’s content address, as well as the file if it is still in the EMC Centera Universal Access file system. EMC Centera is responsible for deleting the files that it stores.

#### Retain files

Certain applications implement the concept of a data retention period whereby an application tolerates a delete or modify request error when the file retention period has not lapsed. EMC Centera Universal Access is able to implement retention periods and pass that information to EMC Centera.

Any request to delete a file that is protected by a retention period will result in the return of a “deletion not permissible” message to the application. In this manner, the combination of EMC Centera and EMC Centera Universal Access can guarantee content authenticity to any application that supports NFS, CIFS, HTTP, or FTP.

#### **Synchronize with EMC Centera**

On a regular basis, usually one minute after the last file update and transparent to the application, EMC Centera Universal Access will synchronize or “write back” new and modified files to EMC Centera. When EMC Centera Universal Access receives each new content address back from EMC Centera, it stores the content address locally and marks the file as synchronized. Because EMC Centera communicates with EMC Centera Universal Access rather than with the application, the content address is managed by EMC Centera Universal Access.

#### **Flush file data and retain file metadata**

EMC Centera Universal Access uses policies to decide when to replace objects in its local cache to keep disk utilization below 70 percent. As the disk becomes more populated, an algorithm monitoring the capacity of the disk cache determines if it is approaching 70 percent full. If so, EMC Centera Universal Access uses an algorithm to determine which files to invalidate (i.e., delete).

Even though EMC Centera Universal Access has deleted the file data from the disk, it maintains the file metadata and continues to present the file as available. When the application requests a file that is no longer stored in cache, Universal Access will read the object from EMC Centera and provide it to the application. Universal Access will retain the file on its disk cache for future reads until the cache reaches 70 percent full and then the invalidate process is run again.

#### **Back up and restore metadata**

EMC Centera Universal Access automatically backs itself up to EMC Centera on a nightly basis. EMC Centera Universal Access captures its state, including the metadata map of content addresses to file names and all relevant system settings and stores this information on EMC Centera. The content address of the nightly backup is automatically e-mailed to configurable addresses for safekeeping. The restore capability is designed to restore a Universal Access system to its prior state on the same or different server with a simple set of commands.

#### **Manually fail-over to Standby Universal Access and Recover**

Enhanced Availability mode maximizes the availability of EMC Centera-resident data and minimizes data loss in the event of a motherboard or network interface card failure on an EMC Centera Universal Access system. Enhanced Availability employs two EMC Centera Universal Access servers working together in an Active and Standby relationship. In Enhanced Availability mode, all files are written to and read from an Active EMC Centera Universal Access, and all subsequent interactions with EMC Centera are managed by the Active EMC Centera Universal Access. At the same time, the Active EMC Centera Universal Access is periodically (every five minutes or 500 updates, by default) publishing detailed Activity Logs (containing complete metadata maps) to the Standby EMC Centera Universal Access. In the event the Active EMC Centera Universal Access fails, the Standby EMC Centera Universal Access is made active by manual intervention and immediately queries EMC Centera for all data written since the timestamp on the most recent Activity Log. Upon receipt of that data, an up-to-date metadata map is created and the Standby EMC Centera Universal Access takes over as the Active EMC Centera Universal Access.

## **Implementing EMC Centera Universal Access**

### **1. Environment verification and hardware acquisition**

EMC Centera Universal Access is a software application that runs on a standard Intel server platform with standard LAN connections.

There are two hardware platforms that support EMC Centera Universal Access software: the Dell 2950 Phase III and the EMC Centera generation 4 and 4 LP nodes. These hardware platforms differ in their architecture and implementation characteristics.

**For the Dell hardware platform, the customer supplies:**

- A Dell PowerEdge 2950 Phase III with Dual-Core 1.6 GHz Xeon processors, 2 GB RAM, and 6x300 GB SAS drives. The Dell PowerEdge 2950 Phase III has onboard gigabit NICs with copper connections, and dual power supplies. The customer orders the Dell hardware directly from EMC. Dell supports the hardware.
- A gigabit LAN switch to connect EMC Centera, Dell 2950 Phase III, and application servers.

**For the EMC Centera generation 4 and 4 LP node hardware platform, the customer supplies:**

- Generation 4 or 4 LP EMC Centera node hardware with an Intel Pentium processor, 1 GB RAM, and four 500 GB, 750 GB, or 1 TB drives. The EMC Centera node has an onboard gigabit NIC with copper connections. The customer orders the EMC Centera node hardware from EMC either in groups of four or singly; EMC supports the hardware.
- A gigabit LAN switch to connect EMC Centera, the EMC Centera Universal Access node, and the application servers

## 2. Software installation, configuration, and licensing

**Configuration:**

- Performed by EMC or by the customer for Dell PowerEdge 2950 Phase III hardware platform, performed by EMC Customer Service for EMC Centera node hardware platform.
- Customer must provide qualified hardware platform (e.g., Dell PowerEdge 2950 Phase III or generation 4 or 4 LP EMC Centera node).
- Use customer-purchased CD for software installation, and serial number for licensing purposes.

**Software installation:**

- Power up the hardware
- Connect to local IP network
- Load Universal Access installation software
- Follow prompts for required info (e.g., EMC Centera address, NTP, e-mail destinations) software licensing
- EMC Centera Universal Access software is licensed via the EMC Powerlink® customer portal, using the license key that accompanies the EMC Centera Universal Access software kit.
- Licensing EMC Centera Universal Access software entitles the customer to support from EMC and access to future updates, patches, and documentation.

## Supported environments

The table below identifies the combination of client operating system platform and access protocols that have been tested with EMC Centera Universal Access.

| Platform         | Protocol                                      |
|------------------|---|
| UNIX             | NFS V2, V3 over TCP                           |
| Linux            | NFS V2, V3 over TCP                           |
| Windows          | CIFS  |
| IBM System i     | NFS V2, V3 using Integrated File System (IFS) |
| Internet Clients | FTP, HTTP (HTTP used only for retrieval)      |



EMC Corporation  
Hopkinton  
Massachusetts  
01748-9103  
1-508-435-1000  
In North America 1-866-464-7381  
www.EMC.com

**Take the next step**

For more information on EMC Centera Universal Access, contact your EMC sales representative.