



# Data Storage System for IAP

Since its foundation in 1932 the Irkutsk Aviation Plant has produced over 20 types of aircraft developed by virtually every aircraft design group in the USSR and Russia. These innovative machines have been exported to 37 nations worldwide. In 1997 IAP became the first Russian aircraft manufacturer to achieve compliance with the ISO 9002 international management quality standard.



## OBJECTIVES

Today's aircraft industry could hardly exist without top-of-the-line information technology. IAP's business strategy includes regular upgrading of its hardware and software systems. A major overhaul was planned in order to start the design of MS-21, a new short to mid-range airliner. In particular, it was necessary to deploy a large fault-tolerant multi-tier data storage system that would support fast access to archives. The contract was awarded to Jet Infosystems, who suggested an EMC-based solution.

## SOLUTION

**Aleksandr Kotenko, head of the project**, recounts: *«We started with systems analysis at IAP. Several types of data were identified that differed in disk space required, relevance to current work and frequency of use. We also realized that the existing data storage system was too obsolete to upgrade. Its components, however, could obviously be used for archival storage of less critical data to save earlier investment. EMC experts also suggested developing a mechanism for relevance-based classification of data and continuous replication to protect data integrity».*

Jet Infosystems partnered with EMC experts to develop a fully fault-tolerant distributed system located on two physically separate sites, i.e. at the main and the redundant data center. Hardware at both sites is virtually identical in terms of functionality and performance. Fig.1 shows a schematic diagram of the system.

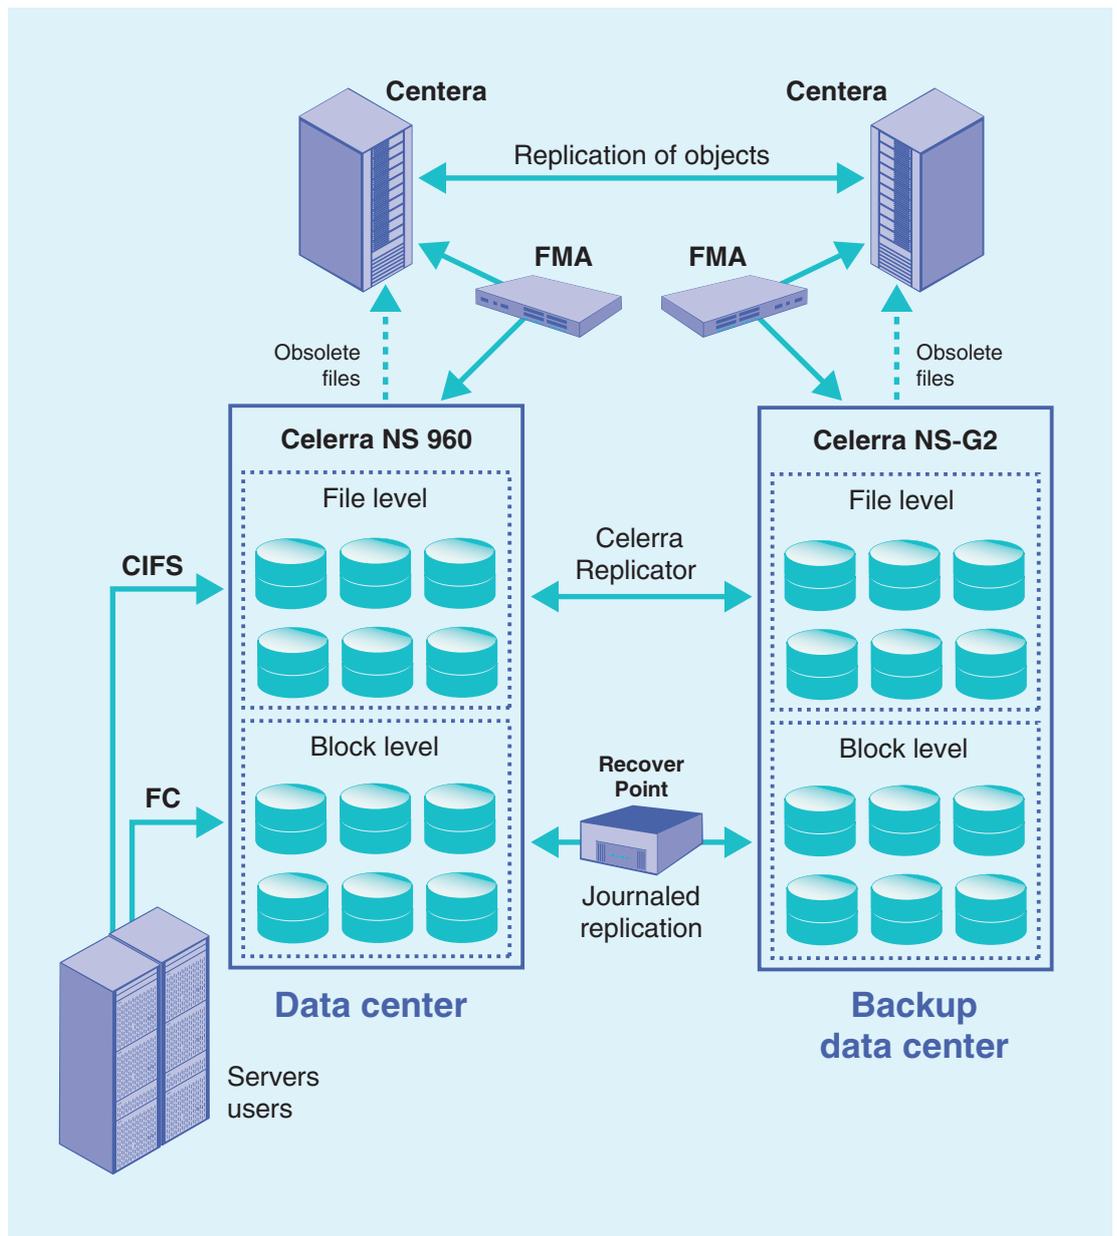
### Current Data Storage

*«We wanted to build a unified multi-tier storage system, and Celerra NS960 proved to be the platform of choice, – Kotenko says. – This array contains several disk types for a variety of purposes».*

The 146 GB 15K FC disks are used to store data that require fast-performing applications (such as database software); the 300 GB 15K FC disks work with CAD systems, store frequently used files and design documentation; the 600 GB 10K FC are designed for file content. The 1000GB 7.2K SATA disks are employed for storing large volumes of consecutive data such as video files, software installation packages and other rarely used large files.

As seen in the diagram, storage resources in the system are divided into two classes for file access and for block access. At the file level users can access data through shared folders. The block level is designed to supply servers with dedicated disk resources for database storage, creation of specialized file systems and the like. The combination of file and block protocols in a single disk array is an evident advantage since it reduces the number of data storage systems used.





### Archival data storage

**Kotenko** goes on to explain: «*Jet Infosystems and EMC have jointly developed a system where information may be stored not only in the Celerra system but also in the EMC Centera archive. The EMC Rainfinity File Management Appliance used for this purpose is capable of transparent data migration between the current data storage facility and the archive. Low-demand files are "retired" to the archive yet restored almost instantly if necessary by using shortcuts*».

Server disk resources are also optimized by means of DiskXtender software that continuously monitors file systems for free space. Once a certain threshold is exceeded, DiskXtender moves files to the EMC Centera archive.

Archival tasks such as saving, data access and removal are managed by means of a special Centera API interface. In addition, the Centera system can automatically deal with data location optimization, self-healing, load balancing, adding capacity and other issues.



## Archive protection

Centera uses CPM mirror copying for data protection. As a file is saved, the application may specify a minimum retention period during which the file cannot be deleted. Information is therefore protected from human error, application crashes or sabotage. Once the retention period expires the file is deleted.

Another level of protection from major disasters is the opportunity of data replication by several Centera systems located on different sites using the CentraStar tool. In this case access to data at the redundant center is very easy, since it would be automatically requested through the Centera API interface if the basic Centera system becomes unavailable.

## Data Replication Between Data Centers

EMC Celerra Replicator V2 and EMC RecoverPoint are used for data replication between the main data center and the redundant center at the file level and the block level respectively. EMC RecoverPoint is a fault-tolerant cluster of servers with pre-installed software; should a server fail, the cluster would continue working. The cluster connects to existing networks and performs continuous journal replication of disk volumes at the block level by taking snapshots to be stored as a journal in a chronological order. The system may thus be restored to any recorded point whatever is the reason of the incident (server failure, data corruption, software malfunctioning, human factor, natural calamities of any other force majeure).

## OUTCOME

The unified multi-tier data storage system is less expensive to operate and significantly more efficient. Due to the flexible settings, policy data may now be prioritized according to access time requirements and frequency of use. Protection level for various data priority classes may also be selected depending on data importance.

**Alexander Shishanov, head of systems division at IAP**, sums up: *«The joint solution by Jet Infosystems and EMC fits our needs perfectly as a reliable high-end system. Data storage space has increased dramatically; even with our growth prospects there is now enough spare disk space for 2 or 3 years. Flexible data management and file classification by relevance reduce our storage costs. Moreover, many routine procedures are now automated».*

*«Multi-tier data storage under this project involves just one disk array with several disk types. This simple yet efficient technology may be used by virtually any business area dealing with large amounts of data processing»*, – says **Roman Volodin, manager of the data storage systems division at Jet Infosystems**.





14 Bolshaya Novodmitrovskaya St.,  
building 1, 127015 Moscow, Russia  
Phone (495) 411-7601  
Fax (495) 411-7602  
E-mail: [info@jet.msk.su](mailto:info@jet.msk.su)  
[www.jet.msk.su](http://www.jet.msk.su)

