sgi®

Global Data Sharing

Use distributed resources, without duplicating data

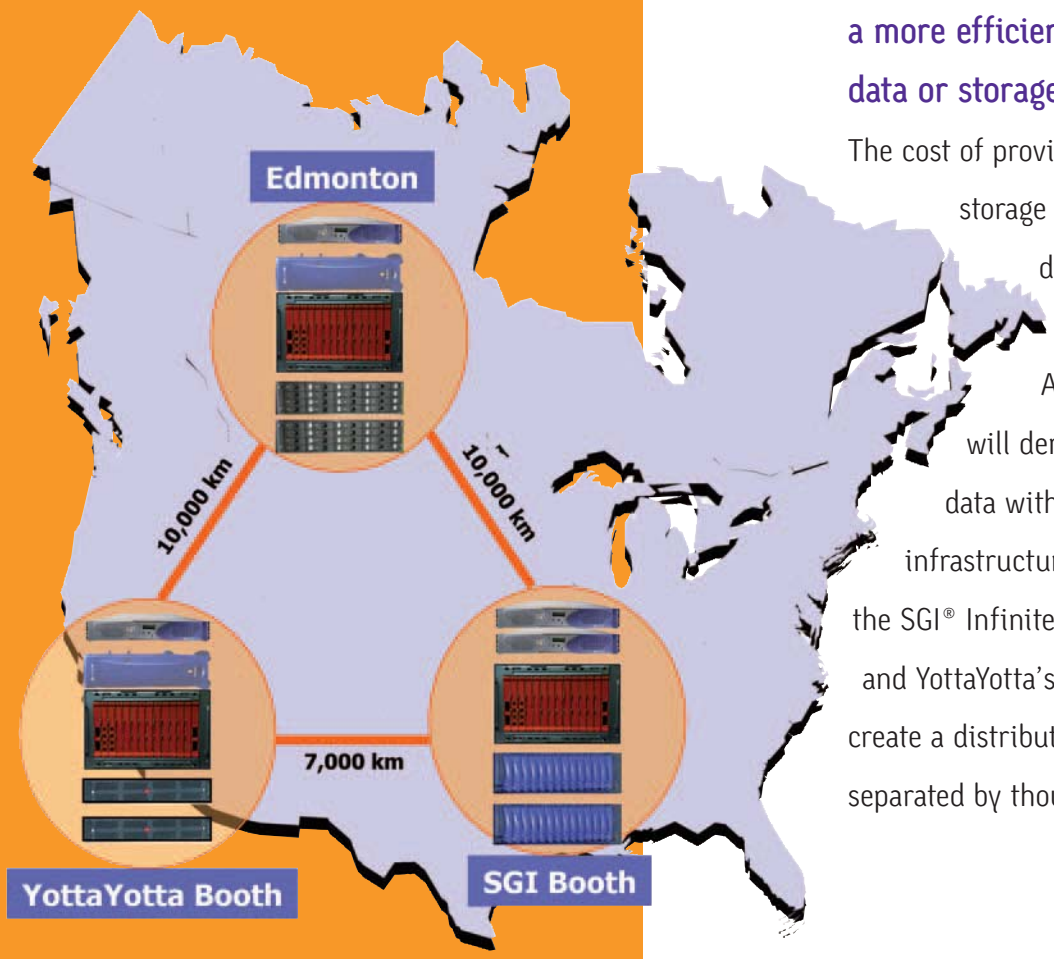

Figure 1: Logical layout for the distributed data centre live demonstration.

Today's businesses must share information in a more efficient manner without duplicating data or storage infrastructure.

The cost of provisioning and managing independent storage archives at multiple geographically distributed sites can be staggering.

At SC2003, SGI Inc., YottaYotta Inc., will demonstrate a new strategy for sharing data without needless duplication of storage infrastructure. The live demonstration deploys the SGI® InfiniteStorage Shared Filesystem CXFS™ and YottaYotta's distributed block system (Y2DBS) to create a distributed data center involving three nodes separated by thousands of kilometres.

Global Data Sharing

Use distributed resources, without duplicating data

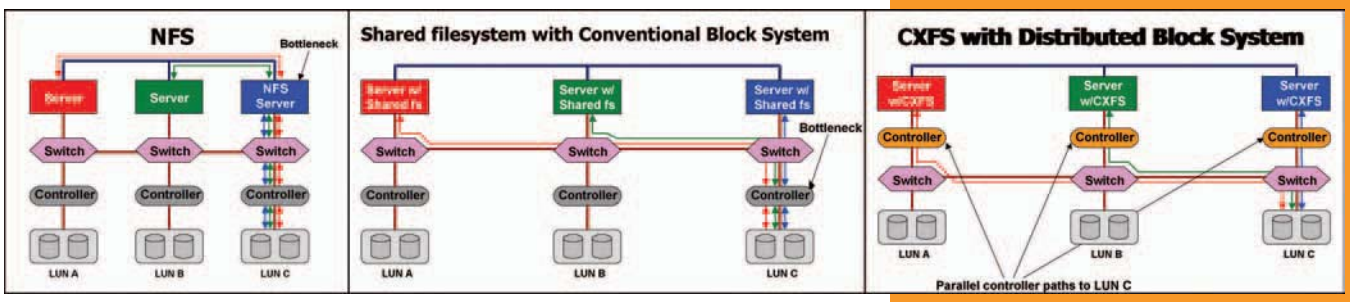

Figure 2: (a) Network File System, (b) Shared Filesystem with Serial Block System, (c) CXFS with YottaYotta Distributed Block System.

Data Localization and Meta-data Localization

Conventional Network File Systems must transmit data to remote servers through a single NFS master server. This master server often seriously restricts I/O. Furthermore, when the data is transferred over larger distances even more serious performance penalties result [See Figure 2].

CXFS removes the data sharing bottlenecks of traditional data sharing mechanisms like NFS and FTP by allowing all the systems on a SAN direct, concurrent access to the same data at the same time – without data mounts or copies. However, data access over large distances are still subject to performance degradation as longer distances introduce latency.

YottaYotta's distributed block system leverages a unique strategy to localize data and meta-data where it is needed when it is needed. By exchanging meta-data (rather than the data itself), performing geographically aware caching, and by migrating metadata directories according to access patterns, YottaYotta ensures that all CXFS systems have peered read/write access to a consistent data image without having to physically transport data unless it is needed. Also, YottaYotta parallelizes data access across an array of storage controllers to ensure that, when data must be transferred across the WAN, it does so at wire speed.

Performance Results from live SGI and YottaYotta Network Tests

In a recent trial between the Canadian cities of Edmonton and Ottawa—separated by 4,700 km—SGI and YottaYotta demonstrated the power of the combined CXFS and Y2DBS solution. Figure 3 shows CXFS performance with and without Y2DBS during the live field trial. For further information about this trial and the CXFS and Y2DBS distributed data access system, please contact:

YottaYotta – *www.yottayotta.com*
SGI Professional Services – *www.sgi.com/services/*

The SGI and YottaYotta solution provides:

- Improved data access and dramatically reduced WAN traffic
- Transparent yet secure access to remote data
- Near-local I/O performance to a globally consistent data image regardless of distance
- Near linear scalability of I/O to shared data regardless of distance
- Flexibility to centrally archive data from multiple sites
- Heterogeneous data sharing

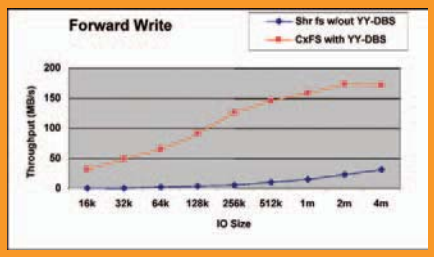

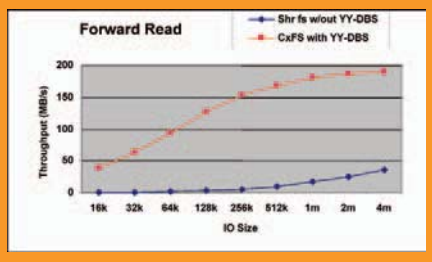

Figure 3: Comparison of CXFS I/O performance with and without Y2DBS.

[11.17.2003] 3621 J14445