Econnectix Corporation
14 North Peoria Street, Suite 2H
Chicago, IL 60607
www.econnectix.com

Table of Contents

1 Introduction

Managing software costs 7x more than the hardware to run it. IT storage costs are doubling at staggering rates, representing up to seventy-five cents of every dollar spent on information technology. In addition to cost, IT infrastructure is steadily expanding requirements not just for increasing capacity, but also for expanding functionality. Storage and accessibility of millions or even billions of connected objects across a network is becoming a necessity. To address and support the current and emerging requirements will require a new class of database designed around the networking implications of the new infrastructures.

Enter **XPress XSN** for Storage Networks – the association database designed to help build scalable networked storage solutions that can track millions of objects, mobile or static, and manage distributed storage metadata at the highest available performance levels and availability.

2 Service Domains

Helping to build a new paradigm for tomorrow's storage networks, XPress XSN addresses the needs and requirements for massive, efficient storage of identity, mobility, locality, presence, and content, essentially the metadata needed to manage distributed storage networks.

The network storage market is emerging at a pace that cannot be ignored. Already, it is a multi-billion dollar market that forecasts on-going remarkable growth; the adoption of Microsoft's .NET and other similar web services reinforces this prediction.

On the one hand, data storage is growing exponentially – today's storage solutions cannot store tomorrow's data requirements. *On the other hand*, how applications locate and use data is changing. Today's data management solutions cannot account for the diverse ways that tomorrow's applications will locate and use data. XPress XSN is the association database designed to manage the underlying key-value associations that manage the indexing and metadata that these applications need when distributed in the machine room or across the network.

To support these new applications will require a new kind of database, above and beyond the **relational database** model, **directory service**, and retasked **DNS**. The most important information describing distribution and availability of data must be managed in real-time. Storage solutions will require tracking and management of information characterizing:

- **Identity**. How objects are determined to be unique.
- **Mobility**. Data moves and applications need to follow it.
- **Locality**. As data moves, different technologies can be used to access it efficiently.
- **Presence**. As data moves, it needs to be discovered.

To support identity, mobility, locality, and presence, storage networks must provide:

- **Performance** networking for metadata management of static *and* mobile data.
- **Scalability** without performance depreciation heretofore unimagined.
- **Network** synchronization across sites for data and metadata redundancy.

Current database technology can be deployed to address these, but was not designed to do so; and it was not designed to provide equivalent computational and financial economies when so tasked. XPress XSN addresses exactly this gap.

XPress XSN design does not address the generality of the relational database model, does not use the replica distribution model of directory services, or the DNS hierarchy, although it can do some, parts, or support components of all of these. Alternatively, XPress XSN provides networked performance unimaginable to relational design, global distribution not suited for directory services, and real-time update capacity unavailable to DNS architecture. XPress XSN isolates a specific function: mapping and unmapping relationships that go on and off the network wire, at new levels of **performance**, **scalability**, and **availability**. For storage applications wanting to manage metadata efficiently in the machine room or in distribution across the wire, XPress XSN has no peer competitor.

3 Service Goals

Storage networks are not the only benefactor, but will be one of the early benefactors of this new database technology. As storage networks expand in scale, complexity, and distribution, they require the performance, scalability, and networking efficiency of an association database designed to manage the objects and metadata describing the objects stored within, outside, and among storage installations.

XPress XSN provides high-availability data storage and retrieval for the crucial indexing and metadata supporting Content Delivery Networks (**CDN**), Storage Area Networks (**SAN**), and Network Attached Storage (**NAS**) – of any scale – at the cost of network latency. Indeed XPress XSN can create arbitrarily large single- or multi-site object stores that can resolve every item's location on a network in near-real-time. As a network optimized database for storage networking solutions, XPress XSN offers specific advantages:

- Up to 2-3 orders of magnitude (100x to 1000x) greater performance than relational databases
- Linear scalability to support any capacity need
- Atomic data consistency between memory and disk without compromising performance advantage, even when distributed across the network
- Faster *replicated* service than non-replicated service in other technologies
- Fast recoverability
- Real-time record synchronization on the network
- Real-time network-wide data consistency
- Geographic distribution limited only by network latency.

Providing such functionality permits constructing flexible, dynamic, distributed policy management.

4 Service Needs

Out of this growth scenario, specific needs become increasingly critical to the adoption and growth of storage networks. What needs does XPress XSN address for network storage applications?

1. **Performance**. The need to locate, manage, and retrieve data objects quickly and transparently among storage devices across a network;
2. **Scalability**. The need to smoothly scale according to growth, as the number of data objects increases exponentially;

3. **Availability**. The need to handle data distribution.

Traditional databases for storage networks do not scale to meet these needs, and their price-performance does not meaningfully address OEM market needs. Finally, other solutions do not provide efficient, integrated networking embedded into the core of the database services like XPress XSN. XPress XSN not only meets these needs, but exceeds them in a cost-effective form-factor that can smoothly scale to any capacity requirement. Let us briefly evaluate these needs.

4.1 Performance

Storage network service will increasingly need to locate, manage, and retrieve data objects quickly and transparently among storage devices across a network. Today large storage installations do not provide global itemization within the installation: databases are not fast enough to support dynamic data mobility in near-real-time at the key-value level. Worse – traditional databases are not designed to support the *networking* required to provide such service in sufficiently efficient form-factor. XPress XSN can provide per-item tracking in real-time as a multi-purpose index server; it is designed and written to provide manipulation of relatively short messages on networks, such as file metadata or file locality information. XPress XSN permits tracking every networked object in real-time.

Protocols and implementations exist for transport, but increased distribution will require a new logical and functional service to manage the volume of objects and potential for randomized mobility across the networks in which they exist and move. Supporting such service requires not only efficient transport mechanisms, but similarly efficient indexing technology that smoothly can handle the volume and dynamism of objects that move in the new storage networks as they are deployed, grow, change, and evolve. Providing such service will require finding, adding, deleting, or updating data objects quickly, scalably, and with replication for availability. A network-designed association database is an appropriate logical unit to this solution, separating indexing and metadata from content so as to provide transparent awareness of the storage network at any moment in time – efficiently, scalably, and robustly.

XPress XSN was designed and written specifically to address these needs. The result? XPress XSN operates at speeds of at least two or three orders of magnitude faster than other database technologies. **Figure 1: XPress XSN Relative Performance** shows the relative performance differentiation between XPress XSN and other database technologies.

Figure 1: XPress XSN Relative Performance

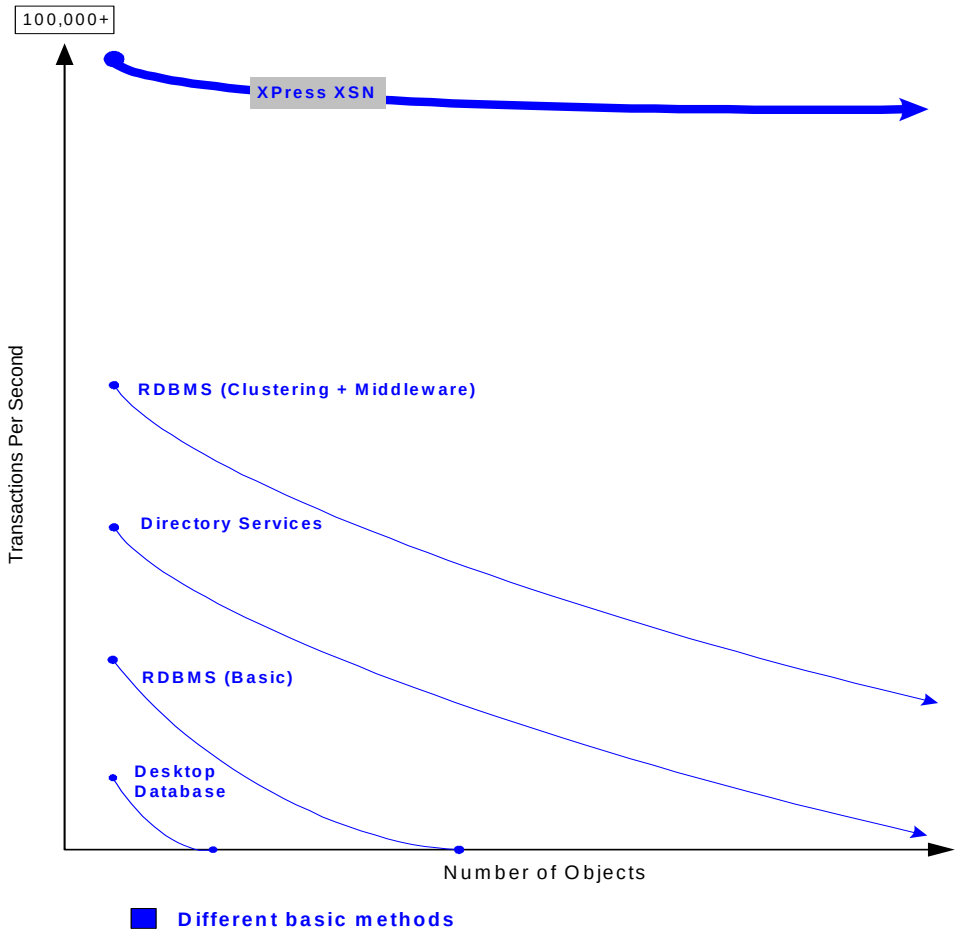

4.2 Scalability

Networked storage service will increasingly need to smoothly scale according to growth, as the number of data objects increases exponentially. Addressing today's high performance management of object locality and metadata management, storage solutions continue to face the well-known trade-offs forced by previous database design. High performance **in memory relational databases** trade disk reliability against speed, and contain processing overhead and generality that is wrongly tasked and poorly designed to address distributed processing across the network. Contrary to "in memory" speed without persistent storage, XPress XSN provides orders of magnitude superior performance than the fastest in-memory relational databases, by removing unneeded design for unneeded function. XPress XSN not only provides this performance, but also provides persistent storage for reliability without performance degradation. XPress XSN provides substantial performance improvements compared against "in-memory" databases while *still providing guaranteed network, memory, disk synchronization* at the sustainable performance levels indicated.

Alternative to relational databases, **directory services** such as LDAP, achieve low latency and high distribution by increasing the number of replicas – but at the cost of update speed. Such design is also not tasked for right function, for the exploding proliferation and mobility of

objects on networks. The greater the distribution of replicas, the longer it takes to update them, even if they can be synchronized. In contrast, XPress XSN achieves distribution through efficient networking, and per-record replication provides high availability at orders of magnitude performance improvement compared against not only directory services, but against all other alternatives.

With the XPress patent-pending clustering technology, XPress XSN applications have no upper bound performance restrictions except network capacity. XPress XSN's core indexing capability smoothly scales to support virtually any capacity of data volume, transaction rate, and client load. Unparalleled performance levels are showcased in the ability to process over 20 billion transactions per day on simple, Pentium-based hardware, or 200,000+ transactions per second, on a single processor. As illustrated in **Figure 2: XPress XSN Scalability**, XPress XSN exhibits linear scalability between the number of transactions per second and the number of servers in a cluster. This linear behavior allows XPress XSN to support systems needing multi-millions of objects and their associated values – *appropriate for managing distributed metadata and file location information that storage solutions require*. XPress XSN does this without the overkill of relational databases pushed to their limits, directory services tasked for managing updates, or the legacy baggage of retasking DNS hierarchy for alternative storage functionality.

Figure 2: XPress XSN Scalability

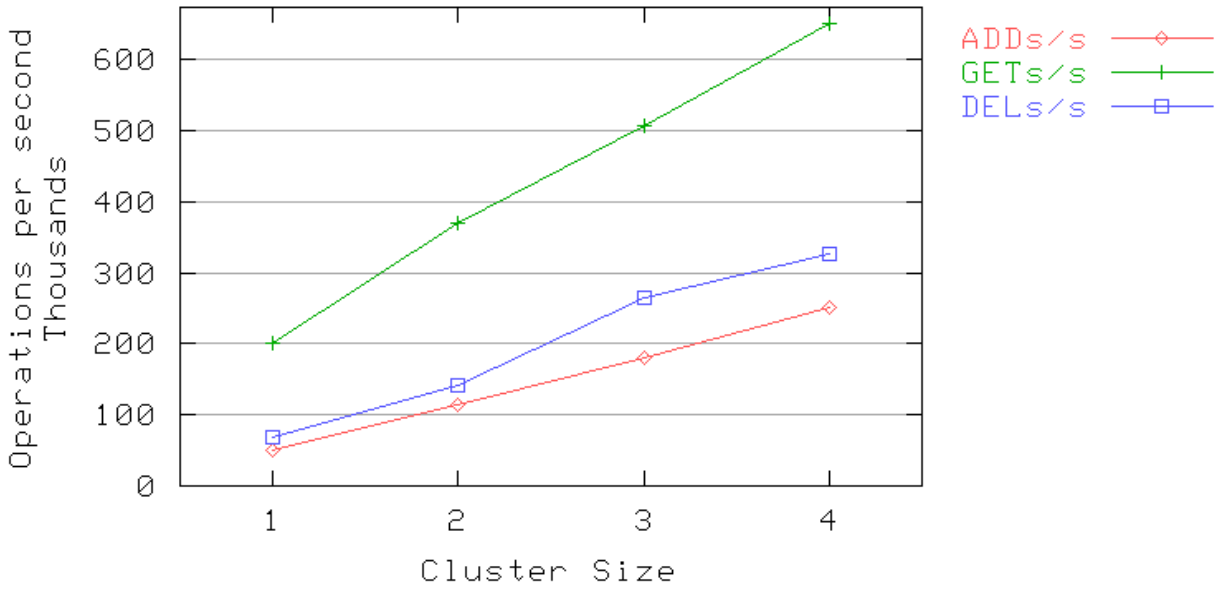

Figure 2: XPress XSN Scalability shows a cluster of 4 Pentium III servers performing 650,000 lookups/second, or 39 million lookups/minute. These tests include round-trip TCP/IP times, running on standard 100mbps Ethernet.

4.3 Availability

Networked storage service will need ever increasing availability of critical metadata as storage becomes distributed. Enterprise data storage solutions are scaling to performance requirements by locating data in multiple locations, sometimes called "the edge". If the metadata is not also multiply-located, the aggregate system performance can degrade. Proper solutions include either:

1. Moving metadata close to the data and/or clients for low latency
2. Separating metadata at the cost of network latency.

A properly implemented solution would minimize the computational and networking overhead to provide either service or some combination of both. XPress XSN allows storage networks to push object tracking and metadata management to the edge or the center, or combinations of both. As data moves, due to changes in usage patterns or organizational structure, the metadata should be able to follow – or lead.

Figure 3 : XPress XSN High Availability

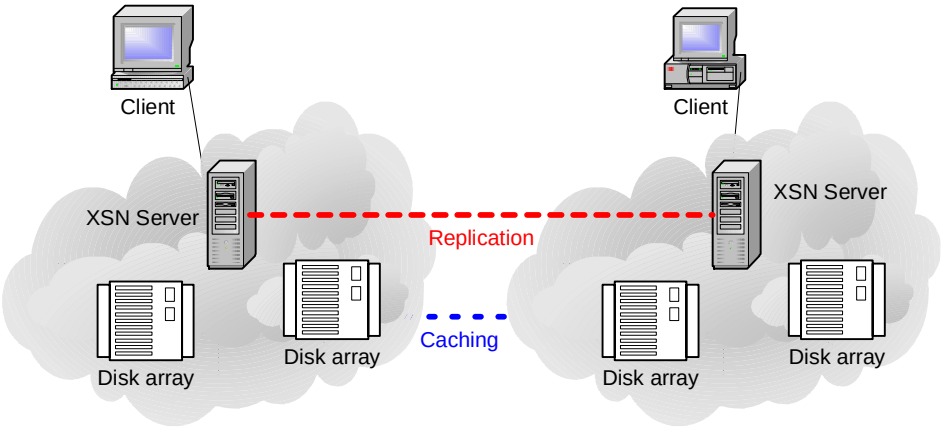

When files are accessed frequently, they can be moved to closer network caches to reduce latency both for awareness of the networked objects, and also for optimal data transfer time. XPress XSN provides unparalleled throughput performance directly servicing fault-tolerant distribution so metadata can be distributed and synchronized in real-time – *at wire speed*.

5 Conclusion

Storage networks will continue to require increasing levels of performance, scalability, and availability. Providing this will require database services designed for object locality and metadata management. Econnectix XPress XSN provides such services using a patented disk, memory, and network synchronization architecture that provides breakthrough performance, scalability, and availability for managing small data elements in networked contexts.