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

Table of Contents

1 Introduction

Small messaging requirements are increasing with the convergence of services crossing the voice/data barrier, and with wide distribution and high availability required to provide low latency services across the network. Telecommunications service providers do not merely need more performance at refined price-points, but convenient scalability for transaction growth, in a replicated high-availability environment, with an altogether different computational-cost basis. These needs can be all reduced to an infrastructure's ability to locate network-wide objects with performance, transparency, scalability, and resiliency. XPress XTN is designed to provide the generic key-value management that short messaging applications require: call management, phone lookups, service routing, service availability, 911, E911, presence and mobility management, instant messaging, and so on.

Providers no longer have to install general purpose databases with painful costs, or special-purpose solutions for specifically bounded problems. XPress XTN provides high performance, convenient scalability, resilient distribution, and a new computational and financial cost-structure unavailable with other technologies for generic background database needs.

2 Service Domains

Helping to build a new paradigm for tomorrow's telecommunications networks, XPress XTN addresses the needs for efficient storage and retrieval in small messaging contexts, such as identity, mobility, locality, presence, and delivery of content in distributed or network contexts.

On the one hand, small messaging applications are growing rapidly: some carriers already deliver 100s of millions of instant messages on an hourly basis; call routing continues to increase; IP convergence continues; number portability has been mandated and will require new infrastructure components; and numbers and types of services are proliferating. *On the other hand*, service providers are forced to exploit available technologies to address current needs with technologies designed for different application purposes, such as relational databases, directory services, and retasked DNS (for example, with ENUM) for common back-end service delivery. Today's data management solutions cannot account for the diverse ways that tomorrow's applications will manage the underlying key-value associations that facilitate the indexing, metadata, and messaging needs that the high volume transaction applications will need in the machine room or across the network.

A new kind of database is needed to *generically* support these new applications. It must go above and beyond the relational database model, directory service, and retasked DNS. It must, in real time, manage information describing distribution, data availability, and delivery. Arriving applications will require information tracking and management characterized by:

- **Identity** - how objects are determined to be unique
- **Mobility** - that data will commonly move
- **Locality** - that diverse applications will need to connect to mobile data
- **Presence** - that dynamic service discovery and data availability must be characterized.

To support identity, mobility, locality, and presence, converged communications networks must minimally provide:

- **Performance** networking to describe static *and* mobile data.
- **Scalability** without performance depreciation, in computationally and financially appropriate form-factor.
- **Network-Memory-Disk** synchronization across sites for data and redundancy at the atomic level.

Current database technology can be deployed to address such needs, but pushes much low-level functionality to the developer, such as replication and distribution provisioning. Current database technology is known not to scale arbitrarily and was not designed to provide equivalent computational and financial economies when forced to scale.

XPress XTN addresses exactly this gap between conventional technologies stretching their design purposes, and the database need to support small messaging manipulation across the network wire – as a generic service for diverse applications.

XPress XTN 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 XTN provides networked performance unimaginable to relational design, global distribution not suited for directory services, and real-time update capacity unavailable to DNS architecture. XPress XTN 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 communications applications wanting to manage high performance needs efficiently in the machine room or in distribution across the wire, XPress XTN has no peer competitor.

3 Service Goals

Communications networks are not the only benefactor, but will be one of the early benefactors of this new database technology. As communications networks expand in scale, complexity, and distribution, they require the performance, scalability, and networking efficiency of a database designed to manage the indexing; metadata describing the objects stored within, outside, and among communications installations; and message delivery for diverse application needs. XPress XTN can create arbitrarily large single- or multi-site object stores that can resolve every item's key-value relationships on a network in near-real-time.

As a network optimized database for communications networking, XPress XTN offers specific technical design advantages:

- Up 100x to 1000x greater performance than relational databases
- Linear scalability to support any capacity need (compute, storage, distribution)
- 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 by network latency.

Providing such functionality lets service providers construct flexible, dynamic, distributed policy management heretofore unavailable.

4 Service Needs

Out of this growth scenario, specific needs become increasingly critical to the adoption, proliferation, and growth of new services on communications networks. What needs does XPress XTN address for telecommunications infrastructure applications?

1. **Performance**. The need to locate, manage, and retrieve data objects quickly and transparently among devices across a network.
2. **Scalability**. The need to smoothly scale according to growth, as the number of objects and devices increase.
3. **Availability**. The need to handle data distribution.

Traditional databases for communications 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 XTN. XPress XTN not only meets these needs, but exceeds them in a cost-effective form-factor that can smoothly scale to any capacity requirement. We briefly evaluate these needs.

4.1 Performance

Today global itemization can occur through services and embedded technologies tasked for specific purpose: but database back-ends are not fast enough to support dynamic data mobility in near-real-time at the key-value level at high volumes of randomized data exchange, such as appropriate for number portability. Worse – traditional databases are not designed to support the *networking* required to provide such service in an appropriately efficient form-factor. XPress XTN can provide generic per-item tracking in real-time as a database designed and written to manipulate small messaging on networks, such as call routing information, GPS location, SMS messages, ESNs, Instant Messages. XPress XTN permits tracking individuated networked objects 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 the objects exist and move. Supporting such service requires not only efficient transport mechanisms, but similarly efficient indexing technology that smoothly handles the volume and dynamism of objects that move through telecommunications infrastructures as the infrastructures are deployed, grow, change, and evolve. Providing such service will require finding, adding, deleting, or updating communications objects quickly, scalably, and with replication for availability. A network-designed association database is an appropriate logical unit to this solution.

XPress XTN was designed to address messaging needs efficiently, scalably, and robustly. XPress XTN operates at speeds up to orders of magnitude faster than other database technologies. **Figure 1: XPress XTN Relative Performance** shows the relative performance differentiation between XPress XTN and other database technologies.

Figure 1: XPress XTN Relative Performance

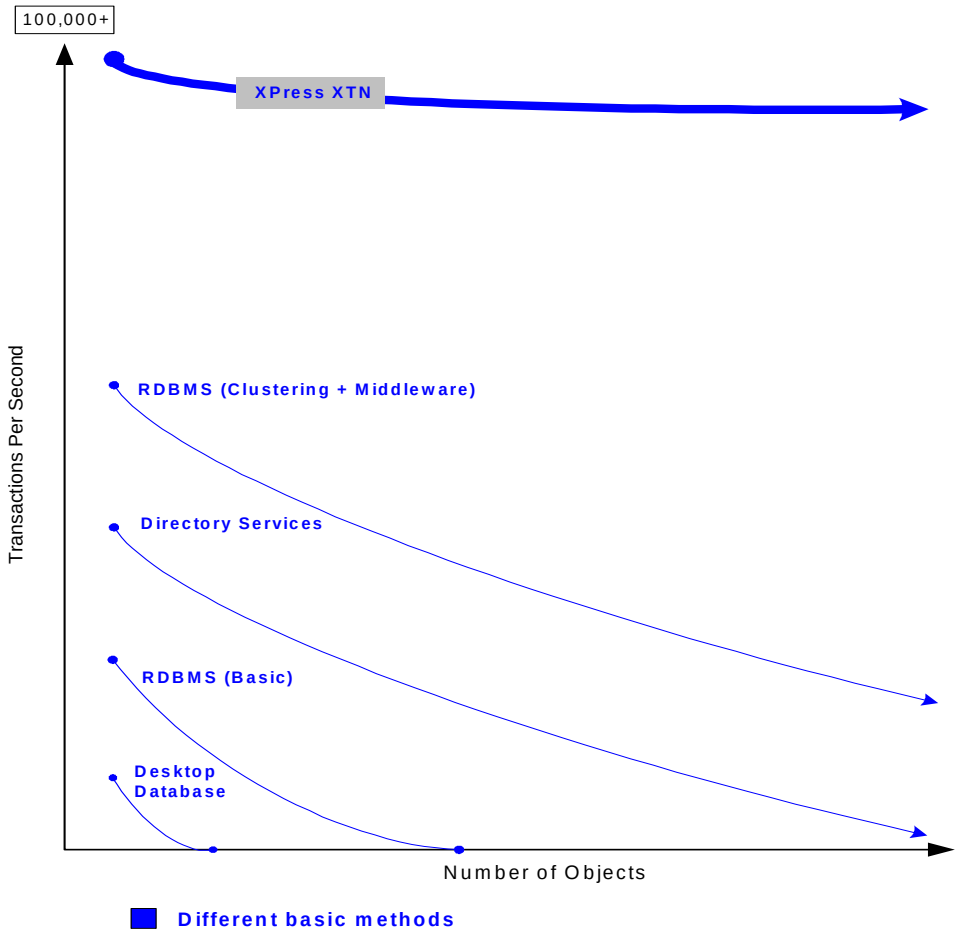

4.2 Scalability

With the proliferation and introduction of new services, smoothly scaling according to growth already is a common design concern, and this will increase as the number of data objects increases exponentially. 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 XTN provides up to orders of magnitude superior performance than the fastest in-memory relational databases, by removing unneeded design for unneeded function. XPress XTN includes *guaranteed network-memory-disk synchronization* at the sustainable performance levels previously 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 the exploding proliferation and mobility of objects on networks. In contrast, XPress XTN 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 XTN applications have no upper bound performance restrictions except network capacity. Unparalleled performance levels are showcased in the ability to process over 20 billion transactions per day on simple, Pentium-based hardware, or 200,000+ peak performance transactions per second, on a single processor. As illustrated in **Figure 2: XPress XTN Scalability**, XPress XTN exhibits linear scalability between the number of transactions per second and the number of servers in a cluster. This linear behavior allows XPress XTN to support systems needing multi-millions of objects and their associated values – *appropriate for managing distributed service management, locality, presence, and availability that communications solutions require*. XPress XTN does this without the overkill of relational databases pushed to their limits, directory services wrongly tasked for managing updates, or the legacy concerns of retasking DNS hierarchy for alternative functionality than it was designed for.

Figure 2: XPress XTN Scalability

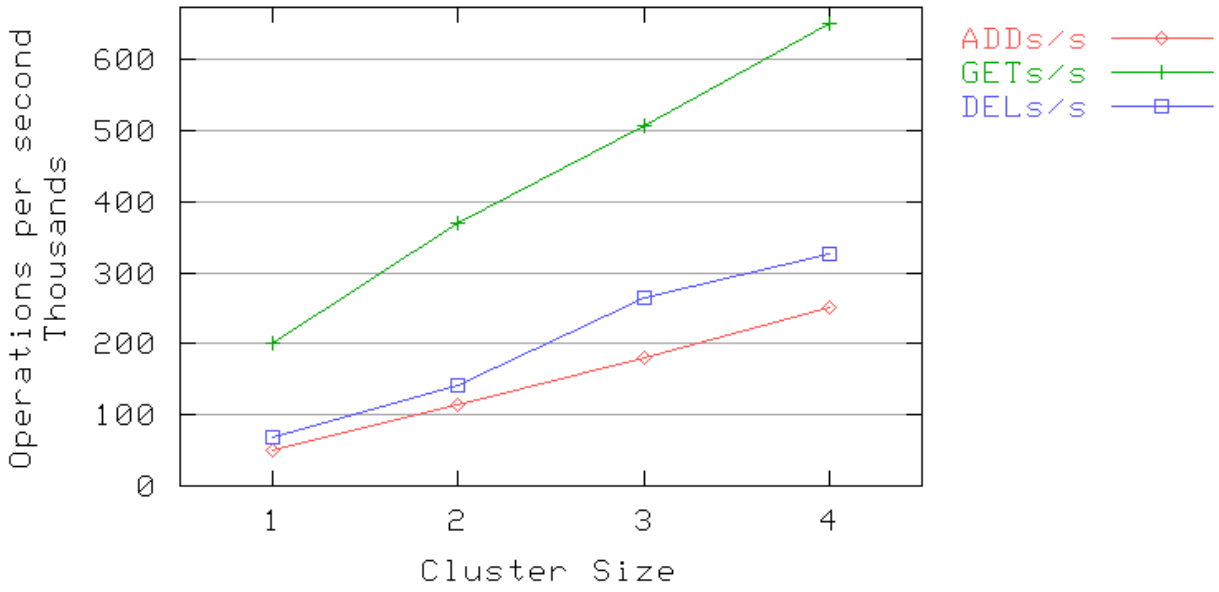

Figure 2: XPress XTN Scalability shows a cluster of 4 Pentium III servers performing 650,000 lookups/second, or 39 million lookups/minute. These tests include round-trip TCP/IP times, running on standard 100mbps Ethernet.

4.3 Availability

Telecommunications services will need increasing availability of critical object information as it becomes distributed. If the object data is not multiply-located, and queries distributed, aggregate system performance can degrade. Conventional solutions typically include either:

1. Moving communications objects closer to clients for low latency. But distributing data increases update time costs.
2. Centralizing communications objects to increase performance through hardware and middleware. But centralization increases network latency costs, which can be unsatisfactory.

A properly implemented solution would minimize the computational and networking overhead to provide either service architecture or some combination of both, while still providing cost-

appropriate scalability. With the economies that XPress XTN provides, applications may push object tracking to the edge or the center, or combinations of both, while still providing reliable and predictable performance.

Figure 3 : XPress XTN High Availability

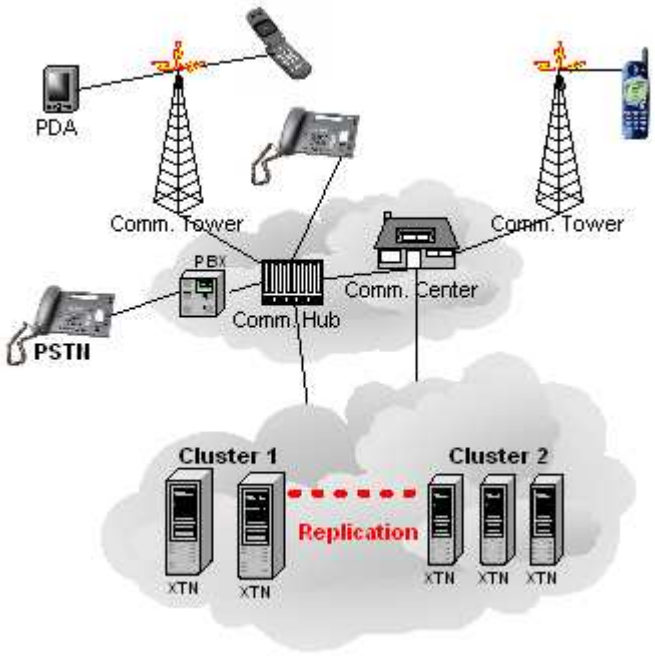

When files are accessed frequently, they can be moved to closer network caches to reduce latency both for awareness of the networked objects, and for optimal data transfer time.

5 Conclusion

Telecommunications networks will continue to require increasing levels of performance, scalability, and availability. Providing this will require database services designed for object management in networked environments. Econnectix XPress XTN provides such services using a patent-pending disk-memory-network synchronization architecture that provides breakthrough performance, scalability, and availability for managing small data elements in networked contexts.