



XPress Association Database

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

Table of Contents

1	Save Today, Scale Tomorrow	1
2	The Net-worked Database.....	1
3	What XPress Delivers	2
3.1	Performance.....	2
3.2	Scalability	4
3.3	Availability	5
4	XPress Advanced Features.....	6
5	Who & What Benefits?	7
6	Conclusion	7

1 Save Today, Scale Tomorrow

Database licensing costs represent a sizeable and direct bite into profits for high performance, high availability applications that anticipate large database images distributed in the machine room or across the network. Moreover, distribution and availability required to provide reliable low latency services across the network are not design foci for conventional database technologies. Accessibility of millions or even billions of connected objects across the network is becoming real, and the need to span multiple high-performance services with a common backend database technology has already arrived. Cost-efficient backend database service is crucial to provisioning these services, but conventional technologies create unfavorable performance-cost ratios, and *real-time updates regardless of scale are not possible* with conventional solutions. Relational database management systems (**RDBMS**) provide generality but are expensive to scale and do not distribute easily and cost-efficiently. Much less, they are designed to compute arbitrary relations rather than to provide efficient object management and small messaging manipulation for network-focused applications. **Directory services** provide distribution and improved cost-bases, but do not provide efficient update or replication performance. Finally, while **DNS** has demonstrated global service distribution and robustness, it was never designed to manage mobile objects, or multiple services concurrently.

Yet a universe of proliferating objects and services continues to expand, and needs for object mobility approach ever more specifically in domains including telecommunications, networking, registry services, security, and storage. Efficient manipulation of key-value associations represents a large portion of the mission critical services underlying these application domains.

Service providers within these broad application domains will not merely need more performance at refined price-points, but convenient and modular scalability for transaction and capacity growth, in replicated high-availability environments, with an altogether different computational-cost basis. These needs can reduce to an infrastructure's ability to locate network-wide objects with performance, transparency, scalability, and resiliency. The XPress Association Database technology is designed to provide the generic key-value association management that applications require when distributed across the machine room or across the network. XPress provides performance, linear scalability, resilient distribution, and a new computational and financial cost-structure unavailable in any other available technology.

2 The Net-worked Database

A new database technology is needed to *generically* support current and arriving high-performance application domains that are distributed across the machine room or across the network. It will not replace the proven utility of conventional technologies for their appropriate-use scenarios, but a new network-oriented database technology must go above and beyond the relational database model, directory service, and retasked DNS for the needs of object manipulation and small messaging in networked contexts. In real time, it must manage information describing distribution, data availability, and delivery. Arriving applications will require tracking and management of objects characterized minimally by:

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

To support identity, mobility, locality, and presence, these new application domains must minimally provide:

- **Performance** networking to describe static *and* mobile data
- **Scalability** without performance depreciation, in computationally and financially appropriate form-factor
- **Availability** of data across the machine room or network

Current database technology can be deployed to address many needs to varying degrees, but frequently pushes low-level functionality to the developer, such as replication and distribution provisioning. Conventional database technology was not designed to provide equivalent computational and financial economies when required to scale like XPress is designed to scale. For XPress, efficient performance, scalability, and availability are design foci.

The XPress Association Database addresses exactly this gap between conventional technologies stretching their design purposes, and the database needs required to support object management and small messaging manipulation across the network wire. XPress provides a generic service for any application domain manipulating centralized or distributed key-value association sets of any complexity.

The XPress Association Database design does not address the generality of the relational database model, does not use the replica distribution model of directory services for performance, or the DNS hierarchy for distribution, although it can do some, parts, or support components of all of these.

XPress Association Database technology provides networked performance unavailable to relational design, global distribution not suited for directory services, and real-time update capacity that DNS hierarchy cannot provide.

For applications wanting to manage high performance needs efficiently across the machine room or across the network wire, XPress has no peer competitor.

3 What XPress Delivers

What needs does the XPress Association Database technology address for object management and small messaging?

1. **Performance** – the need to locate, manage, and retrieve data objects quickly and transparently among devices across a network, at tunable performance levels
2. **Scalability** – the need to smoothly and modularly scale according to growth, to any conceivable degree, as the number of objects, devices, and services increase
3. **Availability** – the need reliably to handle multi-machine or multi-site data distribution.

Conventional database technologies do not competitively meet these needs when measured against XPress. And the price-performance of alternative technologies does not appropriately address OEM market needs when compared against XPress. Finally, other solutions do not provide efficient, integrated networking embedded into the core of the database services. The XPress Association Database not only meets these needs, but exceeds them in a cost-effective form-factor. We briefly evaluate each of these needs.

3.1 Performance


Today global itemization can occur through services and embedded technologies tasked for specific purpose: but database backends are not fast enough to support dynamic data mobility

in near-real-time at high volumes of randomized data exchange, such as appropriate for number portability. Worse – traditional databases are not designed generically to support the *networking* required to provide such service in an appropriately efficient form-factor. XPress is designed and written to manipulate small messaging on networks, such as IP addresses, hostnames, URL management, GPS location, ESN mappings, Instant Messages, etc. XPress permits tracking individuated networked objects in real-time in a generic package.

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 infrastructures as services are deployed, grow, change, and evolve. Providing such service will require finding, adding, deleting, or updating objects quickly, scalably, and with replication for availability. A network-designed association database is an appropriate logical unit to this solution.

The XPress Association Database was designed to address these object management and messaging needs head-on, like no other available database technology, and provides up to orders of magnitude greater performance than other database technologies. **Figure 1: XPress Relative Performance** shows the relative performance differentiation between XPress and other database technologies.

Figure 1: XPress Relative Performance



3.2 Scalability

With the proliferation and introduction of new services, smoothly and modularly scaling according to growth already is a common design concern, and this will increase as the number of data objects increases exponentially. With XPress, developers can choose cost-effective hardware, and then modularly expand that hardware according to growth, rather than having to preselect a more expensive and capacious chassis size at the start. 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 machine room or network. The XPress Association Database includes *guaranteed network-memory-disk synchronization* at the sustainable performance levels indicated in **Figure 1: XPress Relative Performance**.

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. Increasing replicas improves lookup performance, but also decreases update performance. Such design is not tasked for the exploding proliferation and *mobility* of objects on networks. In contrast, the XPress Association Database technology achieves increased performance through clustering which improves both lookup and update performance uniformly – rather than through replica distribution. In addition, XPress achieves distribution through efficient networking, and per-record replication that 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 applications have no upper bound performance restrictions except network capacity. Unparalleled performance levels are represented 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 Scalability**, XPress exhibits linear scalability between the number of transactions per second and the number of servers in a cluster. This linear behavior allows XPress to support systems needing multi-millions of objects and their associated values – *appropriate for managing DNS, distributed service management, number translation services, etc., in a cost-efficient, high-availability platform*. XPress 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 Scalability




Figure 2: XPress 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.


3.3 Availability

Critical object information requires replication when distributing the database, to protect against downtime. Moreover, distributing the object data distributes query load and can improve aggregate system performance by avoiding “hotspot” performance degradation. Conventional solutions typically include either:

1. Moving objects closer to clients for low latency. But distributing data increases update time costs. This design often characterizes the directory solution.
2. Centralizing object management to increase performance through hardware and middleware. But centralization increases network latency costs, which can be unsatisfactory. This design often characterizes relational database solutions.

A properly implemented solution would (a) minimize computational and networking overhead to provide efficient distributed *or* centralized management, (b) provide superior computational efficiencies to drive costs down and performance up, and (c) provide strict modularity to support linear scalability. With the economies provided by XPress, applications may push object tracking to the edge or the center, or combinations of both, while still providing reliable, predictable, and scalable performance. Indeed, XPress allows applications to modularly tune system performance to any desirable degree.

Figure 3 : XPress High Availability



When objects 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. Or object management can be centralized in machine room clusters. XPress permits choosing either or combinations of both topologies.

4 XPress Advanced Features

Telecommunications, networking, registry services, security, and storage infrastructure providers are not the only beneficiaries, but will be the early beneficiaries of this new, network-oriented database technology. As services expand in scale, complexity, distribution, and diversity, they will require improved performance, scalability, and networking efficiency for crucial backend object indexing; metadata manipulation describing the objects stored within, outside, and among provider installations; and message delivery for diverse application needs. The XPress Association Database can create arbitrarily large single- or multi-site object stores that can resolve every item's key-value relationship on a network in near-real-time. Such performance advantages lead to constructing previously unbuildable applications, radical cost reduction for current backend services, or both.

As a network optimized database for object management and small messaging, XPress offers specific technical design advantages:

- Up to 100–1000x greater performance than other technologies, such as RDBMSs
- Linear scalability to support *any* capacity need (compute, storage, distribution)
- Faster *replicated* service than competitor technology's non-replicated, "native" service
- Fast recoverability
- Network-Memory-Disk synchronization for atomic data consistency across service (multi-machine, multi-cluster, multi-site)
- Real-time record synchronization on the network
- Real-time network-wide data consistency
- Geographic distribution limited by network latency.

Providing such functionality facilitates the construction of flexible, dynamic, distributed policy management heretofore unavailable. XPress can provide a common backend that can concurrently support an arbitrary number of diverse services.

5 Who & What Benefits?

Helping to lower costs and to provide a generic backend database solution optimized for the network, the XPress Association Database technology directly addresses the needs of current and arriving application domains. Such service domains include, but are not limited to:

- Messaging
- E911 / 911
- Location Based Services
- Presence and Availability
- Number Translation
- Telecommunications Number Portability
- DNS
- Next Generation Networking
- Directory Services
- Topic Maps
- Access and Service Coordination
- Security & Identification Applications
- Packet and Optical Applications
- VoIP
- Storage management

On the one hand, service providers are increasing types of services provided. Consider how rapidly telecommunications providers introduce new services to increase value to their customers. Or consider even DNS service providers who extend native expectations of service delivery to corporate customers, such as providing guaranteed update time for private-network DNS, or dynamic number translation services for telecommunications carriers. Types and volume of services will continue to proliferate.

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 (e.g., Content Delivery Networks) for common backend service delivery. Yet today's data management solutions were not designed for networked ways that tomorrow's applications will manage their underlying key-value associations.

A well-designed network-oriented database technology does address the object management and messaging requirements that volume transaction applications can use today across the machine room or tomorrow across the network.

6 Conclusion

Distributed object management and small messaging environments continue to require increasing levels of performance, scalability, and availability, and new feature sets designed for networked contexts. The Econnectix XPress Association Database technology provides such service using a patent-pending disk-memory-network synchronization architecture that provides breakthrough performance, scalability, and availability for managing small data elements in networked contexts. XPress design and performance economies immediately lower the bottom-line today, while providing arbitrary scalability for tomorrow's needs.