

ZOTECA™

PUTTING RAPS™ INTO THE NETWORK:
RELIABILITY, ACCESSIBILITY, PRIVACY, SECURITY

Zoteca Data Sharing (ZDS)

Zoteca and RAPS are service marks of Zoteca, Inc.

ZOTECA™



2472 Broadway / Suite 195 New York, NY 10025

Tel: 917-496-6240

Fax: 212-439-4178

www.zoteca.com

1 Data Sharing in Information Technology

1.1 The Problem

In a separate whitepaper on the Zoteca BackEnd (ZBE) we noted that over the past decade, accelerated by the growth of the Internet, information technology (IT) had a new mission: to enhance communication and the sharing of corporate information throughout the organization and across organizational boundaries. As discussed in that whitepaper, the IT department faces enormous challenges in carrying out this mission because of the enormous complexities involved in creating event-driven, asynchronous applications in a distributed network environment. ZBE is an IT tool meant to reduce that complexity.

However, a new layer of complexity is added when IT needs to take into consideration issues of efficiency and safety in sharing data. These issues are not relevant for all distributed data applications. As a rule of thumb, the way to recognize such applications, is to think about corporate systems that are traditionally heavily paper form based.

Paper is extremely *reliable*. When you fill out a paper form, there is no “save” operation necessary in order to ensure that when you finish filling out the form, the data will be there. When properly stored, these documents are always available for reference. The term “paper trail” is highly apt — while not necessarily easy, it is certainly possible to audit paper documents to find where errors might have crept in, or who did what and when. That’s why paper shredding is usually a sign of a cover-up. Someone is attempting to prevent an audit of operations and data reliably stored on paper.

Paper is also very *accessible*, as long as its properly filed. Go to the proper drawer, proper folder, there it is. Need a copy, just go to the copy machine. Need to send it to someone else, FedEx or fax does quite a good job of ensuring that it gets to the proper person. Paper is highly portable, so it is easy to carry with you when you need it for reference.

There are also some nice safety features associated with paper. You can physically limit access so only authorized personnel can access paper documents.

Of course, in large volumes, paper becomes very unwieldy. Storing and searching paper documents becomes an unwieldy, almost impossible task in anything but the smallest organization. Computer automation was rightly welcomed as a savior in corporate information systems. The computer’s ability to store and search through vast amount of data is the only way a modern organization can operate.

The introduction of networks, both local- and wide-area also seemed to promise added efficiency in the sharing of data within and without organization, perhaps eliminating paper forever. But somehow, that hasn’t seemed to pan out yet.

Firstly, computer data is inherently unreliable. One aspect of reliability that computer-based systems are worse than paper-based system is the audit trail. It is very difficult to trace changes in data in computer-based information systems. Another aspect is lost data. Anyone who has ever used a computer has experienced losing data because an application crashes, or a file disappears. Organizations try to minimize data loss through ever-more sophisticated backup and storage technologies. However, these all work based on the assumption that the data is stored in centralized servers. This leads to the second problem of computer data, namely accessibility — access to data is always fastest when it physically resides on the local device. This is true even in local area networks, and even more so in wide area networks (WAN). The problem gets even worse when WAN access is through dial-up or wireless connections. For this reason data users tend to bring the data they need to the periphery or edge of the network. In the recent WTC disaster, centralized corporate databases were quickly backed up. But large amounts of important data, sitting on people’s personal workstations, was irrecoverably

lost. Hence there is an inherent tension between reliability (centralized data) and accessibility (edge data).

Issues of reliability and accessibility relate to the efficiency of data use and data sharing. There are two more important issues of data safety — security and privacy. While these two are related, they are not the same thing. Security breaches can compromise privacy, but privacy issues exist even where security is not an issue. As in all systems, IT systems face the inherent contradiction between safety and efficiency. The more you have of one, the less you have of the other.

Within organizational boundaries, the RAPS issues are not always so severe. But once you get to information systems that require sharing often-changing data across organizational boundaries (the equivalent of paper-based information systems) the RAPS issues become extremely complex.

1.2 Existing Solutions

There are many existing IT techniques for sharing data in these complex inter-organizational environments — email, VPNs, websites, secure FTP or Peer2Peer systems — but all have serious shortcomings.

To begin with, email and P2P systems address the accessibility issues by bringing data to the edge. But they are inherently unreliable, ad-hoc and uncontrolled. In addition, they are difficult to secure.

On the surface, VPNs seem like better solutions since they make use of corporate reliability services, and provide some level of security. However, data needs to be shared between different organizations, across firewalls. Breaking the firewall boundaries means a loss of security the firewalls were supposed to create in the first place. VPNs basically giving full access to the internal network to an external user. For organizations sharing data with outsiders this is simply unacceptable. Secure FTP, while giving more limited access to data, shares all the shortcomings of VPNs. Both are bad in terms of accessibility. If either the client or server get disconnected, there is no way for users to access data. This is especially acute when using modems and wireless devices.

Web-based systems share the worst-features of all these solution on all four dimensions of RAPS.

Another problem to be dealt with is data notification. Email is the only solution that inherently notifies the user that data has arrived. All the others must provide some sort of additional act of notification (another email?) that the data being shared is now available.

In addition, for those solutions other than VPNs that do provide safety measures through user authentication, the users cannot use their standard usernames and passwords from their own organizations' network — they need to remember a new password.

In other words, sharing data is *hard*.

2 The Zoteca Solution

Zoteca's data sharing (ZDS) solves these problems, using a distributed data sharing and storage system. Basically the system works by providing a shared database, the ZStorage. The ZStorage is divided into different namespaces, or partitions, each one having it's own access control. The ZStorage is transactional, ensuring data integrity. In addition to these rather standard features, there are a number of things that make the ZStorage unique:

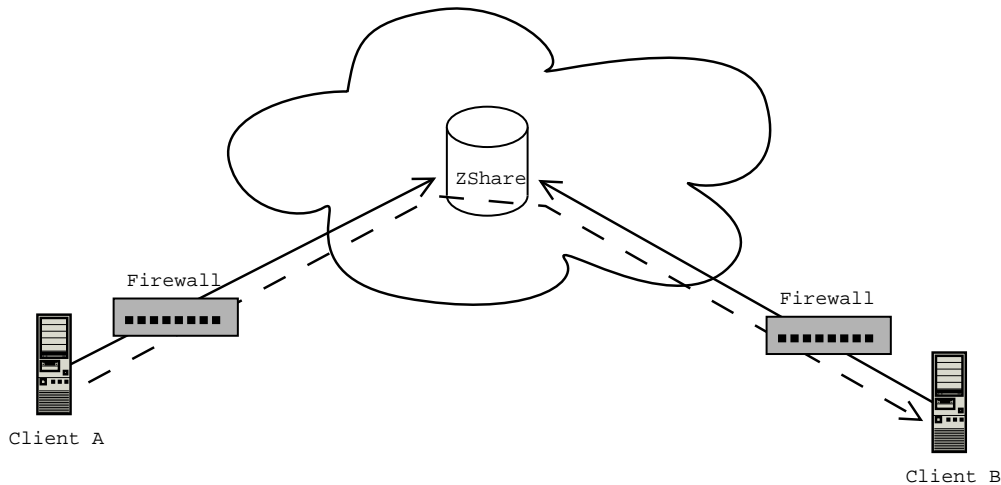


Figure 1: In an example of how data is shared, both client A and client B are connected to the ZStorage. Client A inserts some data into namespace N. Since client B is subscribed to N, it gets notified that new data was inserted, and it automatically downloads the new data. Now this data is available to B even later, when it gets disconnected from the server. Notice that both clients are behind firewalls, and not directly connected, yet the data flows from A to B.

- Data is never deleted from the ZStorage. Once added, data can never be changed, deleted or tampered with in any way.
- Users are notified when data they are subscribed to is inserted. This data notification supports disconnection, letting users be notified of inserts done when they were disconnected, and can be done by clients behind firewalls.

Each namespace in the ZStorage has read and write permissions, which can be granted to specific clients. Using sophisticated cryptographic techniques, these permissions can be checked by proxy servers even when disconnected from the main server.

The ZStorage is built on top of existing databases, RDBMS or OODBMS, but because it is an abstraction layer above them, it allows advanced functionality such as replication across different databases. So, for example, the main server may be Oracle replicated with DB2, whereas the caching proxies use PostgreSQL. The ZStorage itself is not a relational database — it can use an object database as a backend just as well.

3 Benefits of Zoteca Data Sharing (ZDS)

3.1 Reliability

The fact that data is immutable allows using simpler, less fragile algorithms for designing replicated or distributed servers. Creating fail-over and backup servers is also much easier, due to the append-only nature of the data transformations. Since replication begins at the local client, even local changes are

preserved, similar to paper-based systems. Hence ZDS can virtually guarantee that data is never lost or destroyed, and provide disaster-recovery even of edge data..

3.2 Audit Log and Historical Data

As a result of the data immutability in the ZStorage, new data never replaces the old. Instead, new versions of the data are added. The result is a log of all the changes made to the data. This is useful both as an audit log, as well as for viewing historical versions of data for other purposes — a mistakenly deleted file may be recoverable from a version from the last 15 minutes, comparing different versions of data, etc.

3.3 Data Consistency

Since the database is transactional, data moves from one consistent state to another. Data immutability allows changes downloaded to clients to be done in a transactional manner as well, since the client can know whether or not it has downloaded all of the changes to the data.

3.4 Caching and Disconnected Operation

Another benefit of immutable data is that cached data never needs to be verified. In traditional distributed databases, cached data is always suspect, since it may have changed in a different server. In the ZStorage however, data is immutable, and thus can be cached indefinitely. This means that data can be stored as close to the client as possible, and the client can access the local copy.

In addition, this enables disconnected reads — the data a client is reading may all be in the local cache, and thus there is no need to even access the server. The ability to authenticate users even when disconnected makes this ability a reality for proxy caches as well. Disconnected writes using a queue may also be possible, depending on the type of data being stored.

3.5 Data Migration

When a user inserts data into the ZStorage, the subscribers to this entry or its namespace get notified of the data's creation, whenever they next connect. They can download this data automatically. The result is that data migrates to the clients that need it — there is no need to manually go and check for new data, data will arrive where it is supposed to as soon as possible, without human intervention.

3.6 Safety

The ZStorage is only used for sharing specific data that needs to be shared. It is not a database where all the data is stored, or a replacement for existing databases. A VPN opens up all of an organization's system and then tries to limit access. In the ZStorage, on the other hand, only data that has been approved for sharing is explicitly pushed out — and the firewall is never opened up to external users. In addition, the data is shared without a point-to-point connection — the data goes source->ZStorage->destination, instead of source->destination.

3.7 Local Access Control

The ZStorage system can be integrated with different organizations' access control and authentication frameworks. Instead of sharing data between clients, data can be shared between application servers. These application servers can then use the local access control (e.g. Windows NT, or Unix) to restrict access to data they got from the ZStorage. At the same time, they do not need to know about whatever system other application servers use.

3.8 IT Benefits

Scalability Yet another benefit of immutable data is that making scalable servers is much easier than with other distributed or replicated databases.

Speed ZStorage's superior caching abilities, provide a large speed benefit in data sharing.

Cross-platform ZDS supports all major operating systems (Windows-based and Unix-based) as well as operating platforms such as Java (and .NET when it becomes more widely available).

3.9 Organizational Benefits

Evolutionary Integration ZDS integrates with the organization's existing software, policies and network infrastructure, and does not require an expensive redesign of the complete IT system. An evolutionary, future-compatible integration is possible, allowing the piece-by-piece modernization of the organization's data management systems without causing upheaval and disruption.

Interoperability: Each organizational unit, and each external organization, has its own custom policies, and in some cases custom software that needs to be interfaced with. ZDS easily integrates with this heterogeneous environment.

4 Sample Applications

4.1 Patient Records in a Hospital

A hospital has many different units, each with their own applications, databases, access control systems, etc. In addition, the hospital must communicate with external organizations — insurance companies, testing labs, external experts, and others. Patient records need to be accessed by all of these, without compromising the hospital's internal and external boundaries.

Using Zoteca's ZStorage product, a data sharing infrastructure can be put in place between the different units and external organizations. This layer shares the appropriate data, making sure that only the correct data reaches the correct place. For example, a request for a medical test goes to the labs, and the insurance company is notified that a test has been ordered for that patient. The data does not go to any other places, and the patient data and test results remain safely secured behind firewalls without VPN holes. The data sharing system integrates with the hospital's existing applications and databases, and lets each unit use its own access control systems.

4.2 Sharing Files in a Transparent File Sharing

A file sharing application allows users to specify a folder on their machine as “shared”. Files are saved from standard applications into that shared folder. Appropriate means are used to distribute keys to authorized users within and without the organization who are designated to receive the shared files. Any file stored in this folder is automatically distributed to those who should get it.

Data is reliable. If a file is lost, historical versions can always be recovered, even without being stored in the centralized corporate servers. Data is always accessed from the local device. Users can browse the folder off their own harddisk, even when disconnected, by simply looking at appropriate folders on their own computing device (workstation or PDA). Safety is preserved. No one sniffing the network can know to whom data is being sent. All they would see are some encrypted packets. Even in a worst-case scenario where a key is compromised, the only data compromised is the specific files shared, not the corporate database.

Founders

Aron Trauring — CEO. Aron has worked nearly 25 years in technical development and management and international sales and marketing for high-tech companies in the US and abroad, including a stint as Director of European Sales at AMDOCS. Seven years ago he co-founded an interactive agency, MAXIMA Multimedia (<http://www.maximam.com>). He also co-founded an Internet B2B ASP two years ago.

Itamar Shtull-Trauring — Chief Technology Architect. Itamar has worked professionally in software technology development for nearly seven years. He studied computer science and mathematics at Tel Aviv University and is the author of several patents. He most recently served as the chief programmer at an Internet startup.

Board of Advisors

Dr. Mahadev Satyanarayanan Professor Satyanarayanan is the Carnegie Group Professor of Computer Science at Carnegie Mellon University. An experimental computer scientist, he has pioneered research in the field of mobile information access. An outcome of this work is the Coda File System, which provides application-transparent support for disconnected and weakly-connected operation. Key ideas from Coda have been incorporated by Microsoft into a forthcoming release of the Windows NT file system. More recently, Satyanarayanan and his research group have been working on application-aware adaptation, a more general approach to mobile information access. This concept is being explored in the context of a new platform, Odyssey.

Dr. David Chaum — Founder and a member of the Board of Directors of DigiCash Inc., a company that has pioneered electronic cash innovations. In the area of cryptography, Dr. Chaum has published over 45 original technical articles, received over 17 US patents, and founded the scientific organization, the International Association for Cryptographic Research (IACR). Concurrently he created and chaired the Smart Card 2000 conferences and several European Union funded industry consortia, including CAFE, which focused on electronic-wallets and the smart cards they hold. He built up a cryptography research group at the Center for Mathematics and Computer Science (CWI) in Amsterdam and during this time also founded DigiCash.

Dr. Mel Horwitsch — Professor of the Management and Chair of the Department of Management at Polytechnic University and founding Director of the Institute for Technology and Enterprise. He is also Visiting Professor at London Business School. Previously he was Professor and Founding Dean of Management at Theseus Institute in Sophia Antipolis, France, serving on the Theseus Board of Directors and Theseus Scientific Advisory Board. He has written extensively on innovation and technology strategy, particularly with reference to such knowledge-intensive sectors as services, information technology, and telecommunications.

Contact Information

Aron Trauring

Direct: +1 (917) 496-6240

Office: +1 (212) 905-3261

Fax: +1 (212) 905-3266

email: aronst@zoteca.com

Web: <http://www.zoteca.com/>

Address:

2472 Broadway, Suite 195

New York, NY 10025