

Requirements for a Secure Embedded Database

By Suren Behari
Product Manager – Team Developer

September 2004

GUPTA™

Table of Contents

Introduction	3
Typical Implementation.....	3
Insurance Claims Adjusters	3
Pharmaceutical Field Trials	3
Corporate Banking Applications	4
Web Business Site.....	5
Conclusions	5
Development and Deployment.....	5
Connected, Open and Interchangeable	5
Small Footprint	6
Self Managing.....	6
Invisible in Operation.....	6
Operational Management	6
Exception Management	6
Fault Notification	7
Remote Diagnosis and Administration	7
Field Upgradeable	7
Security and Risk Management.....	7
Private.....	7
Tamper-proof	8
Conclusion	9

**Information
vulnerability is
not at the
point of
compromise**

Introduction

This paper outlines the key business requirements that GUPTA has identified as being essential for the deployment of databases containing core business information into environments which are not, or cannot be as secure as the traditional "glasshouse" internally managed systems.

This paper begins by describing the sorts of businesses where embedded databases can greatly improve service quality and profitability; as long as the information jewels of the enterprise can be safely accessed in those environments.

Typical Implementation

In this section we discuss some typical businesses, which would look to make use of secure embedded databases and anecdotally profile the attendant risks. In doing this, we begin to understand that the information vulnerability is not at the point of compromise, but in the ramifications to the organization as a consequence of the information systems and data being tampered with or stolen.

Insurance Claims Adjusters

Insurance claims adjusters are most effective in their roles when they're providing rapid, timely response to their customers' claims. Spending time in the office following up on matters is time away from productive contact with clients; and when a person's crop has been hit by a hailstorm or their only mode of transportation has been stolen, a quick turnaround on the part of the adjuster translates directly into quality of life issues for the clients.

Information vulnerability for these people stems from two areas:

- Confidential client information; which can lead to a loss of trust between client and insurance organization (with attendant liabilities following on); and the strong possibility of the loss of business;
- Data sensitivity; given that the adjusters will potentially cut checks there and then, it is critical to ensure that the data is guaranteed safe: neither touched nor tampered. The consequence of "claims hacking" is not in the direct cost where the check is made out for \$90,000 instead of \$10,000; but in the cost of the system auditing that will need to take place once the security breach is inevitably discovered.

Pharmaceutical Field Trials

In the United States, Pharmaceutical field trials are essential for gathering FDA approval; and the process of shepherding a medication or drug through the various stages of approval takes about 7 years. It is standard business practice for these companies to apply for patents for their compounds at the outset of the process. These patents take place immediately and last for 20 years; and because of the length of

time of the field trials, it generally means that the companies can expect only 13 years of patent protection for the inventions or discoveries. Conventional wisdom has it that the administration process costs one million dollars per day.

Medical and drug technology today is at a state almost unimaginable even 30 years ago; however, even with the advances in biological modeling and biotechnology analysis, much of the work is necessarily empirical: start with a compound with known properties; and tweak it to cause it to have new properties; and then discover exactly how those properties manifest themselves in a biological system.

Having well organized information that is accessible to the trial staff is clearly of tremendous direct benefit, in terms of the time and efficiency savings and more importantly, in terms of the medical information that may be vital in determining how best to proceed with a trial.

If the competition were able to gain unauthorized access to the data, then they could significantly reduce the amount of investigation and research they were required to undertake in order to determine the efficacy or usefulness of a type of compound.

However, if this information were to be compromised by an unethical competitor, the validity of the clinical information would be suspect and significant parts of the trials would need to be restarted. This is the nightmare that is feared by the information professionals and security managers in these organizations.

Corporate Banking Applications

Every corporate organization is concerned about ensuring that funds management is effectively carried out; both on an investment basis, but also a day-to-day basis: overnight money market accounts and so on. Couple this with the increasingly global nature of the marketplace and organizations find themselves needing to be adroit and timely at moving money between accounts and currencies.

To serve these needs, a service sector that provides high-value banking applications to corporate organizations has evolved. They provide the organizations with the software that allows them to set up and track mirrors of the accounts which they maintain with the corporation's banking organizations; and then, in a fashion similar to the home banking product Quicken, from Intuit, move the money from account to account and then wire that information back to the banks as instructions.

Companies need to worry about two types of unsustainable risks in this environment:

- What is the risk of somebody breaking into the application and changing a transfer amount for his or her own purposes?
- What costs will be incurred as a result of the large-scale and far-reaching audits that will need to be undertaken in order to uncover the extent of the financial damage? It is in the loss of productivity; and of customer trust that the most significant portion of the damages arises.

Secure, tamper-proof databases are essential

Web Business Site

The Web is a wonderful way to get information to, and from, your customers. The idea that you could hook your core business systems up to the Internet to allow customers to transact business on their own behalf is both exciting and terrifying.

Apart from operational management issues (the Web is 7x24 whereas most internal systems are 8x5) the essential question becomes: who has access to what information, and what can they do with it?

It is apparent then, that it is not feasible to contemplate putting your ERP or Purchasing System on the Web. The approach that most organizations look to take is to provide a staging database; a gatekeeper if you will, that acts as a bridge between the internal organization and the external enterprise.

As such, the information contained within the database is potentially very important and is certainly very sensitive. Thus it is critical to ensure that this information, which is at its most vulnerable when it is being hosted at the entry point to the Internet, is as safe and tamper-proof as possible.

Conclusions

Each of these situations, based on actual customer contacts with GUPTA, show that secure, tamper-proof databases are essential for ensuring that information can be safely delivered to the new environments where people are doing business: more and more it's not behind the firewall, sitting in front of a computer at a desk; but out in the open where the information is susceptible to compromise.

The rest of this paper speaks to the requirements, which a secure embedded database needs in order to provide the requisite levels of development, operational and security support.

Development and Deployment

This section discusses the needs for the development staff that are charged with designing a system that can be successfully deployed in a wide variety of environments; where the only certainty in implementation is that additional applications will be built to take advantage of the data that is available in the database once built.

Connected, Open and Interchangeable

Since no man is an island and no database is alone, it is important that the database be appropriately open to the world: a variety of applications will need access to the information contained within – standards such as SQL, ODBC are well established; and OLE DB is an important strategic direction for many organization's connectivity standards.

Additionally, in situations where transmission security is less important, Type 4 JDBC support for Java applets and applications is also a requirement.

However, not all traffic is from application to database – organizations employ a mixture of OLTP (Online Transaction Processing), OFTP (Off

Line Transaction Processing) and database replication to ensure that the right information is available as required.

Thus, the embedded database solution must ensure that it has replication connectivity for a variety of diverse back ends – other DBMS systems, ERP systems and UNIX and mainframe hosted systems.

Small Footprint

Work expands to fill the time available, and software bloat on people's PC is a prevalent, and ever-present problem. Accordingly, the database solution must have a low disk- and memory- footprint to ensure that the application is as "sociable" as possible. Reasonable disk space requirements are necessary in order to make field upgrade ability on already congested hard disk space a possibility; and in-memory footprint is essential in order to ensure that the deployment of new systems does not force the organization to upgrade the existing hardware to take advantage of these new systems.

Self Managing

By definition, an embedded database runs within an application that will be deployed in a number of different environments: perhaps on every notebook and PDA that is used within the organization. Accordingly the database has to be able to adjust its performance and behavior as the data within it grows and changes in nature.

Also, because individual applications have very different needs, the database engine must provide full support for administrative operations including, but not limited to, file system management, resource management, security and backup/restore.

Invisible in Operation

Operational invisibility is much more than simply an aesthetic preference. Icons and applications take up real estate on the desktop; and running programs tempt the curious and nefarious. Accordingly, the database engine should start up, run and stop silently and unobtrusively.

Operational Management

Once built, the application will need to be managed through its life: through planned and unplanned changes and outages. This section discusses some of those requirements; from the point of view of "what happens if something goes wrong?"

Exception Management

An "exception" is a technical term that refers to some unintended program behavior – divide by zero, program crash or "abend" in mainframe parlance. It is important that the system be able to recognize that such an exception has occurred and deal with it.

Logic exceptions such as divide by zero should be dealt with entirely within the embedded database. However, sometimes, unanticipated program behavior can cause the database engine to crash. The first response of the system needs to be to check whether the database is running; and if not, recover from it.

Fault Notification

Faults may occur on a one-off basis: perhaps the action of another program or application has inadvertently caused the database engine to fail. Unfortunately, this is all too common in the operating systems on which embedded databases are typically deployed Windows 9x and even Windows NT.

However, over time, and with a number of installed systems, patterns of behavior will become apparent and it is important that these faults be made available as part of a store-and forward facility.

Remote Diagnosis and Administration

Periodically, or in response to a failure event, the organization's DBA or MIS people may wish to gain access to the database system for routine management, data changes or further investigation of a problem. Accordingly, the database engine should have remote diagnosis and administration capabilities.

Field Upgradeable

Over the course of time, as product improvements or repairs become available organizations may want to take advantage of them. Sometimes, a whole application suite upgrade is in order; more often, only a single component needs to be changed.

Generally speaking, the application suite upgrade occurs on a regular timed basis, in response to sustained development which is being undertaken on a well-defined schedule. In this event, the entire application needs to be upgraded.

In the case of the embedded database engine, it is important that the product provide features that allow for that component to be upgraded. This provides a "black box" environment, which assures MIS management that the cost of ownership, after deployment, will not be unfeasibly large. So, if a problem were to be uncovered, then just that affected component could be changed.

Security and Risk Management

One of the challenges that face information security professionals is providing a secure, tamper-proof database environment to store this sensitive information. Accordingly, security measures are of paramount importance in the embedded database environment.

Private

If information is a corporate asset, then it is important to the organization that it be protected from prying eyes. Security is a relative goal, not an absolute one; and so the organization, in conjunction with their security policies, must establish the level of encryption that is required to keep the data private.

There are three potential areas of vulnerability, each of which must be strongly guarded:

- On the wire, between the client application and the database. Network analyzers may compromise networks; but no less susceptible is the notebook, where people may

- introduce sniffers or tracing products, which will monitor the conversation between the application and the database.
- In the database: the storage systems (the files) for the database must also be strongly encrypted, because the data must remain safe and private even if the database engine has somehow been shut down.
 - In the backup files and log files; because the information in there contains copies of the information contained within the database.

Although many database vendors have simple pseudo-encryption offerings, these are manifestly inappropriate for safeguarding corporate assets. Accordingly, organizations must evaluate products based on modern cryptographic standards – DES and triple DES and the like.

Worryingly, some of these database vendors also provide log-deciphering tools. In the original environments in which these databases were used (academic institutions or behind the corporate glasshouse) having low-level access tools was appropriate, because the files were protected with multiple layers of security. However, in the mobile- and web-world, no such assurances are possible and no such tools are desirable: if a hacker were able to read your log files and discover the username and password with which you logged on; all of your private information becomes open for their inspection, theft, destruction or modification

Tamper-proof

Finally, there comes the issue of unauthorized intrusion and tampering. It is one thing for someone to steal a copy of data; but how much worse is the threat of someone unbeknownst to you, accessing the data and changing it? What damage does it inflict on the validity of your clinical drug trials; or the effective date of a foreign currency transaction? It is these indirect costs, which typically have the most costly, prolonged, and disruptive effect on your business.

Clearly, it is the responsibility of the database vendors to ensure that the data, once in the database, remains sacrosanct. Tamper-prevention algorithms must be built into the heart of the storage systems to ensure that data cannot be modified, even if the database engine has been shut down or otherwise stopped.

There are a number of alternative solutions which present themselves to the interested reader; and many of them are high-performance offerings which originate from the days of modem-based data communication – CRC (Cyclic Redundancy Check) and so on. However, these offer only a modicum of defense against the determined hacker, who can make use of freely available tools to alter the database file and then alter the signature CRC.

Accordingly, a database engine must offer a variety of alternatives: from low-grade CRC to the ultra-secure, one-way mechanism of SHA (Secure Hash Algorithm) which offers a one-way cryptographic standard digital “signature” for each object that it is protecting.

It is vital to ensure that any tamper-prevention solution addresses the conflicting needs of database performance and granularity of protection. Page based solutions are the most appropriate, since every database

engine, no matter what its lock- and resource-management architecture, reads and writes the file storage system in pages; and a page-based tamper prevention mechanism offers the best of both worlds.

Database authority and security play major roles today in all organizations, due to the increasingly sensitive information applications have access to. SQLBase Treasury Editions offer total protection against unauthorized data access. Encryption is offered at the server level, database level, and database page level, resulting in end-to-end security, a database industry first.

Conclusion

It is the goal of many organizations to extend the ways and places in which they do business; however in order to support these needs, high quality, sensitive information must be available.

Thus, it is imperative that organizations arm themselves against attacks, both internal and external, to ensure that the information, the crown jewels of the corporation's asset base, remains safe and secure.

Responsible data administrators, security officers and business process owners must evaluate all aspects of the information management solution; and must require of their database vendors that the embedded databases which they use can address each of the items described in the paper.

Copyright © 2004 Gupta Technologies, LLC. GUPTA, the GUPTA logo, and all GUPTA products are licensed or registered trademarks of Gupta Technologies, LLC. All other products are trademarks or registered trademarks of their respective owners. All rights reserved.