TANEJA GROUP
TECHNOLOGY ANALYSTS

SOLUTION PROFILE

# Solix Secure Test and Development Solution – A New Frontier in Securing Sensitive Information

April 2007

THE NET NET

Sensitive information is increasingly finding its way into the hands of malicious individuals either through external breaches or insider thefts from employees or contractors. Database applications like ERP and CRM house vast repositories of sensitive and private information such as salary and pay grade information, customer data, Social Security numbers, credit card numbers, and financial data. This data is highly regulated both externally and internally, and the risk of exposing this sensitive information can result in stiff penalties and legal liabilities. In fact, in a March 2007 Taneja Group research study of large enterprises, 57% of end users reported that a data security breach would cost over $500,000 in damages to their organization and 31% stated that a breach would cost them over $1M.

Most enterprises maintain strict controls and security over production database applications and their data. However test development environments for these types of applications represent an Achilles heel. End users state that creating Test and Development (Test/Dev) environments for this critical data results in security risks, hefty manual procedures, and heavy storage requirements.

Solix has responded with a solution that securely and cost-effectively automates data cloning in Test/Dev environments. This solution profile will discuss the data security context within firms operating today and will highlight the key challenges that we hear from end users about creating, maintaining, and securing test development environments. Lastly, we spotlight Solix's, Secure Test and Development solution and its key benefits.

## IT Facing a New Reality in Data Security

The business context that firms must operate within has evolved dramatically over the past five years. Data security, data privacy, and compliance issues have moved to the fore as a slew of high profile breaches have showcased the vulnerabilities of companies' sensitive data. IT is being forced to adjust its processes and approaches to match this new business context. The Test/Dev environment is one such area where IT has a strong need to re-assert controls and security to meet compliance and regulatory legislation, to safeguard data from insider tampering, and to prevent external breaches or leaks.

From our research, the following three drivers are causing IT organizations to re-evaluate and re-plan how they create, manage, and maintain their Test/Dev environments:

- **Comply with government and industry regulations.** The very database applications that are the most critical in Test/Dev environments are the same applications that are the most highly regulated. Most enterprises must comply with external regulations including Sarbanes Oxley (SOX), HIPAA, Gramm-Leach-Bliley Act (GLBA), and others. Penalties for non-compliance can be severe.

- **Safeguard data from insider threats.** Corporate environments are rife with internal threats to data. System administrators, database administrators, and application developers all maintain high-level access to critical data, and may leak or expose sensitive data either intentionally or inadvertently. Test/Dev environments must have data protections and security in place that guard against both accidental and deliberate threats to production data.

- **Prevent external breaches.** Outsourcing and nearsourcing are common business practices, but they make it that much harder to protect data and to comply with regulations. Companies remain liable for security breaches when outside contractors are accessing copies of their crucial company data. In this context, securing Test/Dev environments is crucial to avoiding stiff legal liabilities and penalties for non-compliance.

## Today's Challenges of Creating, Maintaining and Securing Test/Dev Environments

Cloning the production systems, including structured and unstructured data, is a fundamental procedure in the Test/Dev process.  In order to ensure high quality testing of applications before they go into productions, QA and application developers need to use realistic production quality data so they can ensure the application will perform to exacting specifications in production.  As a result there is a strong need to clone and copy production stores for use in application development, quality assurance (QA), application testing and training, and user acceptance testing. However, cloning production data sets is increasingly challenging given enterprise data growth rates and the critical and highly regulated and sensitive nature of this data.

**Challenge #1: Labor Intensive.** Manual cloning requires a massive effort spanning a variety of cross-functional stakeholders. Cloning large production data sets impacts storage, network bandwidth, host servers, and databases; therefore involving at the least network administrators, storage administrators, database administrators, system administrators, and business application owners. Time spent cloning data is also problematic given slow file replication technologies like File Transfer Protocol (FTP) or Remote Copy (RCP). FTP securely transfers files between hosts that are using different file systems or character sets such as EBCDIC or ASCII, and Remote Copy (RCP) is the remote version of a copy command. Cloning a 500 GB database using FTP or RCP can take days of data movement and manual management, and Test/Dev often requires multiple clones.

**Challenge #2: Inefficient Storage Consumption.** Since the Test/Dev environments must preserve relationships between data, cloned production database applications are usually full clones consuming vast amounts of space of Tier 1 disk. The problem worsens as extra clones are created for additional or parallel development and testing work. Many organizations that we speak with report maintaining seven or eight individual copies of production data. Therefore, a single 500 GB database can easily consume over 3 TB of Tier 1 storage capacity.

**Challenge #3: Legal Liabilities & Security Risks.** Furthermore, data and user access is usually not as secure in the Test/Dev environment, as extremely sensitive information is suddenly divorced from the careful security measures that usually surround it in the production environment. The results of careless or malicious handling of this critical data can be catastrophic.

For example, enterprise IT departments and DBAs carefully restrict usage of the database data within the network. However, if a payroll processing development project is outsourced, the outsourcer gains immediate and full access to sensitive and highly regulated employee data. Nonetheless, the corporation that outsourced the project is still liable for ensuring the security and protection of that data. In another example, take an internal development project. The production database is carefully stored on primary disk and is only accessible to a few authenticated users. But the clone of the database is unwittingly placed on common networked storage, making it casually available to any member of the development team – or worse, to any network user or consultant who can access that storage.

## The Business Case for a Secure Test/Dev Solution

In spite of heavy overhead and security risks associated with traditional cloning, companies commonly make about eight manual copies of each production database for Test/Dev environments. Depending on the size of the database, manual clone refreshes demand two to eleven days of staff time alone. Automating the cloning process can reduce this time to less than a day, and with block-based replication technology no longer than an hour. Storage resource requirements for automated clones also shrink, with instance subsetting or pointer-based cloning reducing required storage by as much as 50%.

Human and storage resource savings can easily total hundreds of thousands of dollars a year, and with a secure cloning system the savings grow in proportion to lowering non-compliance risk. For example, the staffing costs of cross functional personnel involved in a manual clone can be slashed if the test/dev cloning process can be automated. As a result, valuable IT resources can focus on more strategic projects that directly add value to the business. Moreover, organizations can reduce capital expenditures on storage capacity by replicating the clones from tier 1 to tier 2 storage. Assuming a 50% reduction in the size of a clone through instance subsetting, this can easily result in savings of almost $100,000 in storage costs alone for a 500 GB database that must be replicated 8 times.

From our discussions with users, they report the following key savings and efficiencies in their environments:

1. **Automate to Reduce Errors:** IT administrators in charge of cloning a production database application might replicate the entire database and its associated application files using a copy application like FTP or RCP. However, this consumes large amounts of storage and network bandwidth and takes a large amount of manual oversight.

2. **Use Recent Production Data for Better Accuracy:** Cloning can be such a painful process that Test/Dev environments end up working with outdated copies. Using a cloning solution that quickly and efficiently clones recent production data results in a more accurate data set, ultimately improving business intelligence.

3. **Secure Sensitive Data to Protect Against Legal Liabilities:** Production databases often contain highly sensitive and regulated data, data that unauthorized users should not see. When production database clones travel to outsourcers, trainers, and QA testers, those users should not be able to view any sensitive information.

4. **Leverage Efficient Block Replication to Compress Time to Staging:** Cloning at the file level requires very long copy windows. If the Test/Dev solution can take advantage of block-level replication technology, then the time to clone an environment can be compressed to a fraction of the time that it used to take. A policy-driven automated process will also allow IT to reduce its reliance on software tools and manual scripts, which saves money on purchases, human resources, and error-prone procedures.

5. **Cut CAPEX and OPEX Costs.** A Test/Dev cloning solution cuts cloning time from days to minutes, slashing human resource requirements. The solution should also support replicating cloned data onto less costly Tier 2 storage such as a NetApp filer with SATA drives, resulting in a 50% storage reduction difference between full clones and subsets. In addition, a Test/Dev cloning solution can streamline clones by cloning only the data and data relationships that the Test/Dev environment actually needs.

## Spotlight on Solix

Founded in 2001, Solix Technologies is focused on helping end users solve their compliance requirements and implement Information Lifecycle Management (ILM) policies and procedures throughout their storage infrastructure. Its flagship product, Solix Enterprise Data Management Suite (EDMS), organizes and manages all enterprise data including structured (packaged and custom applications), semi-structured (email) and unstructured (documents and images) data. Solix Secure Test and Development solution is a standalone module built on top of EDMS. It enables the enterprise to securely and efficiently clone non-production databases for Test/Dev environments.

Solix Secure Test and Development solution combines data security techniques to protect individual data elements with the ability to streamline or shrink clones of production databases by removing large parts of the data sets that Test/Dev environments do not need. This streamlining capability – achieved through a process called instance subsetting -- results in compact clones that replicate quickly, use significantly less storage resources, and streamline Test/Dev environments.

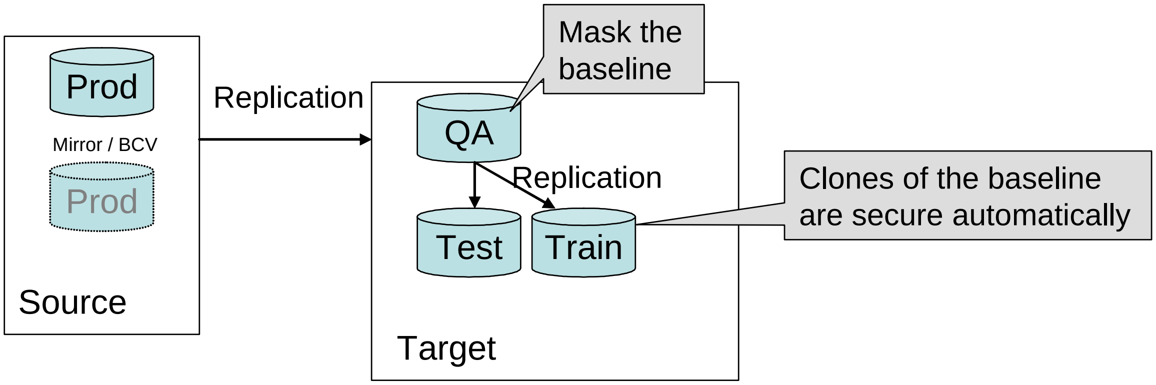

## Efficient Cloning

One of the core storage management operations of enterprise IT is the ability to make multiple usable copies of data sets. This is applicable for a variety of enterprise processes including disaster recovery, testing, analysis, and development environments. However, copy creation technologies are less than ideal in production cloning where the Test/Dev environment requires current data formatting, structure and relationships. Traditional copy creation can yield time gaps in the production copy, impact the production environment with quiescing, require code customization and scripts, and demand manual intervention and management.

Solix uses a technology called instance subsetting to clone subsets for Test/Dev while preserving relational integrity. Subsetting removes large data sets that are not needed in the Test/Dev environment. Removing the data from the clone eliminates the risk of exposure for that data. Another significant benefit of instance subsetting is reduced storage because the clone is significantly smaller than full copies of the production database. Subsetting integrates database and application cloning and is policy-driven.

The compact clones retain a realistic look and feel to the non-production data for accurate Test/Dev purposes, while securely shielding the actual data from the Test/Dev environment workers. The users still have the database structure they need, including application business rules and database relationships, data integrity and accuracy, and proper database access and control including allocating rights to roles.

### Data Protection Techniques

Securing cloned database applications is a fundamental piece of the Test/Dev puzzle. Solix provides a policy engine, multiple security techniques, and knowledge bases to automate the secure cloning process. Since even a single Test/Dev breach of sensitive data can be disastrous, Solix provides a critical piece of technology for protecting companies from data privacy and compliance penalties and liabilities.

Solix Secure Test and Development solution comes with a pre-populated knowledge base containing over twenty populated data security algorithms, and supports the ability to add custom algorithms as needed. IT can define security policies at the application module, transaction, table or column level. The same repository also operates for related operations such as instance subsetting.

Users do not have to be data security experts to choose the correct set of algorithms. Solix's user interface lets users select the sensitive database tables and columns they wish to protect and their security parameters. Solix maintains a comprehensive knowledge base of metadata concerning an enterprise application's data model, which allows it to automatically identify all places where a column is referenced in a given application and apply secured values generated by the optimal data security algorithm.

Solix's knowledge base includes metadata on application data models from major enterprise applications like Oracle, PeopleSoft, SAP, and custom-built applications. This allows Solix to apply data protection techniques to application data while preserving its referential integrity – vital for providing a structurally correct but completely secure clone for the Test/Dev environment.

### Secure Elements

Solix uses scrambling, masking, encryption, nulling out, substitution, and shuffling techniques to secure database clones. All of the techniques employ the metadata repository for applying policy-driven automation and a data model repository to observe enterprise application data models.

- **Data Scrambling** uses an algorithm to scramble data so that it is completely indecipherable. Only when unscrambling algorithms are applied can the data be read again.

- **Data Masking** identifies columns of data and replaces existing characters with designated characters or numbers. The method preserves formatting for testing reports or user interfaces.

- **Encryption/Decryption** encrypts data into special characters and destroys formatting so that the database is unreadable. Decryption keys revert data back into its readable form. Solix uses the DES algorithm.

- **Nulling Out** replaces an entire column with null values. Nulling out is a limited technique as key data columns cannot always be made null, but is a simple and straightforward procedure.

- **Substitution** randomly replaces real production data with realistic-looking fake data and preserves formatting. This substituted data has the same look and feel as the real data for accurate Test/Dev purposes, but does not correlate with the original data or data related to it.
- **Shuffling** is similar to substitution in that it randomly replaces data with fake data. The false values are obtained by shuffling column values, for instance by deliberately mismatching first names with last names.

## Key Benefits of Solix Data Security Solution

### Benefit #1: Streamline non-production databases for Test/Dev environments.

After Solix creates a masked clone of the production database, IT can choose to create streamlined copies of the clone by removing large sets of data that are not necessary for a given Test/Dev task. The remaining subset is smaller, taking up fewer storage resources and resulting in shorter testing times on limited subsets.

### Benefit #2: Easily clone using the most recent production data.

Since cloning with Solix is simple and secure, it is easy to clone updated databases for Test/Dev environments. The compact and secure clones can themselves be cloned without threatening sensitive data or overwhelming storage resources.

### Benefit #3: Leverage replication technology at the block level.

Solix supports enterprise block replication software from major storage vendors like NetApp and EMC, so corporations can leverage their existing replication technologies.

**Solix Deployment Scenarios**

One of Solix's Data Protection characteristics is its flexibility when creating clones and replicating them to different storage tiers. For example:

1. *Tier 1 and Tier 2 storage on different platforms with no SAN.* This environment requires host-based or file-based replication to create an initial clone, such as Veritas Replication Exec or Volume Replicator. Full clones can take time to replicate, but streamlined clones take much less.
2. *Tier 1 and Tier 2 on different platforms with a SAN.* This setup requires block-based replication. Users should leverage data protection software like NetApp's Topio Data Protection Suite (TDPS) or EMC Recover Point (formerly Kashya).
3. *Tier 1 and Tier 2 on the same platform with a SAN.* In this case, the environment should leverage array replication technology such as EMC SAN Copy or NetApp FlexClone to make full clones or snaps.

**Benefit #4: Cut costs.**

Smaller clones result in as much as a 50% storage reduction difference between full clones and subsets. Because Solix supports block replication to secondary storage, the enterprise can leverage existing replication applications and replicate to less costly Tier 2 storage, such as a NetApp filer with SATA drives.

## Taneja Group Opinion

Test/Dev environments create tremendous challenges to the enterprise in the form of serious security risks, storage consumption, and high manual overhead. A single production data breach can easily result in millions of dollars in liability for the company, yet today most companies have not addressed the security of their test and development environments. As a result, they have left themselves open to severe legal liabilities and stiff penalties.

Within this context, Solix offers a clear, hard-cost ROI solution for securing and automating Test/Dev environments. Solix data security techniques are instrumental in protecting Test/Dev database application clones from unauthorized access while retaining full database integrity. Solix can also create streamlined clones, which significantly reduces the overall data footprint of a Test/Dev clone and results in dramatic cost savings. And its automation, knowledge bases, and policy engine dramatically cut OPEX costs by automating manual procedures.

We have already seen significant, high profile data security breaches at large corporations due to lost unencrypted tapes and stolen laptops with sensitive information. Taneja Group believes that it is only a matter of time until a major storage security breach headline will be attributed to an unprotected test and development environment. We urge corporations to reevaluate whether the security policies in place today are sufficient to protect sensitive information in outsourced or test and development environments from falling into the wrong hands. For those with significant outsourcing and insecure test development environments, we urge you to consider the strong ROI potential that Solix provides.

*NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*