

::ISUG

techcast series

# A Real World Case Study: Ensuring Data Privacy in ASE

October 17, 2006

SYBASE®

::ISUG

techcast series

## A Real World Case Study: Ensuring Data Privacy in ASE



**Cindy Bean,**  
Vice President of Events & Media,  
International Sybase User Group

**SYBASE®**

::ISUG

techcast series

## A Real World Case Study: Ensuring Data Privacy in ASE



**Rajnish Chitkara**  
Sr. Development Manager  
ASE Security Team  
Sybase Inc.



**Barbara Banks**  
Staff Software Engineer II  
ASE Security Team  
Sybase, Inc.

**SYBASE**<sup>®</sup>

::ISUG

techcast series

## A Real World Case Study: Ensuring Data Privacy in ASE



**Tad Hawkins**  
Vice President, IT Architecture  
ACS, Human Capital Management Solutions.

SYBASE®

# Agenda

- **Encrypted Columns Feature Description**
- **ACS/HCMS Protects Privacy Using ASE**

::ISUG

techcast series

# Encrypted Columns Feature Description

**Barbara Banks**  
Staff Software Engineer – II, ASE Security Team

**SYBASE®**

## Purpose of encryption - Protection of data privacy

- **What is data privacy?**
  - From a user perspective data privacy is the right to privacy in the collection of sensitive data in digital form to protect against fraud, identity theft or unauthorized use of data
- **Why is it important?**
  - Compliance with legislative mandates, recommendations and expectations
  - Public trust in e-business
    - **Customer bears the risk of fraud**
  - Employee trust in employer

# Encrypted Columns

## Objective

- **Data privacy through encryption**
- **Protect data “at rest”**
  - in the database
  - in backup tapes
  - in replication queues
- **Combined with SSL network encryption, provide end-to-end protection**

## Released with

- **ASE 12.5.3a. Now merged into ASE 12.5.4**
- **ASE 15.0 EC. Now merged into ASE 15.0.1**

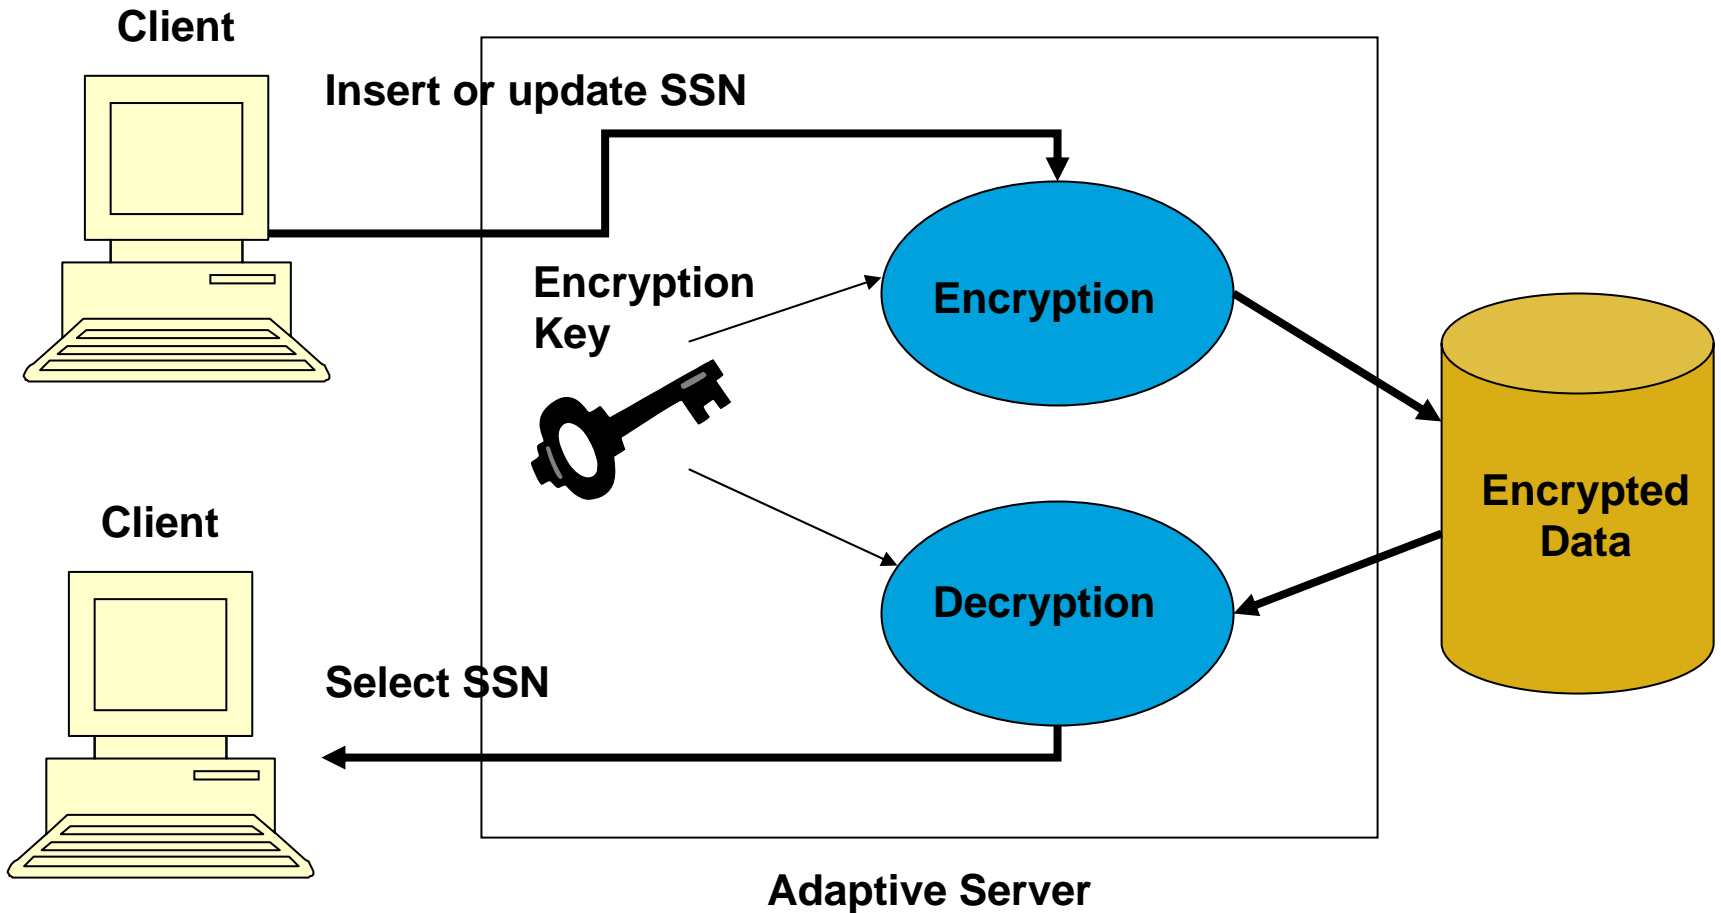
## Functionality

- **Column level encryption**
  - Encrypt only sensitive data
  - Minimize performance and disk space overhead
- **Application transparency**
  - No need to change application, triggers to encrypt and decrypt data
  - DML and SELECT statements continue to operate unchanged
    - **When DECRYPT permission is granted**
- **Performance optimized**
  - ASE optimizer performs efficient searches and joins of encrypted data to minimize the number of decryption operations required

## Functionality contd ...

- **Easy to use**
  - One alter table command to encrypt existing table of data
  - Key creation and management through simple SQL statements
- **Deep integration with ASE and related Sybase products**
  - Supported by RepServer, DDLGen, Migration Tool, ASE Plugin, BCP utility, Power Designer
    - **Replication of data in encrypted form**
    - **Ability to securely migrate keys and data**
- **Separation of roles through decrypt permission**
  - Separation of operator role from privileged users
  - Separation of right to insert data from right to select data

# Encrypted Columns



## Configuration


- 1. Enable encrypted columns**
- 2. Create system encryption password**
- 3. Create encryption key(s)**
  - Stored encrypted in dbname..sysencryptkeys
- 4. Alter table to encrypt data column(s) with named key**
- 5. Grant decryption authorization to users or roles**
- 6. Continue using same old SQL**
  - ASE Automatically encrypts and decrypts column data on INSERT, UPDATE, SELECT; in WHERE clauses

# Encrypted Columns

## Performance Characteristics

- **ASE matches encrypted data based on the ciphertext value, where possible.**
  - Reduces the number of decryption operations per query
  - Allows the use of indexes on encrypted columns

Select \* from emp where ssn = '100-99-1234'



## Performance Characteristics (cont'd)

- **Joins across tables are optimized the same way**
  - Provided the same key is used on joining columns

Select \* from emp e, phone p where e.ssn = p.ssn and phone\_type = 'mobile'

**phone**

ssn	phone_type
0x53b46798301	home
0x9a0ff4658c32	home
0x4277cb02107	mobile

**emp**

ssn	empname
0x53b46798301	rajnish
0x9a0ff4658c32	barbara
0x4277cb02107	tad

# Encrypted Columns

## Performance Characteristics (cont'd)

- **Comparison of logical I/Os between same query with**
  - Index on searched column 'order.o\_c\_id'
  - 1.3 million rows in orders; 700k rows in customer

```
Select c_first, c_last from orders o, customer c
      where o.o_c_id = c.c_id
      and c_last like 'B%'
      and c_first like 'B%'
      order by c_last
```

Query/data characteristics	Number of I/Os
Plaintext data	11486
Encrypted data with ciphertext matching	11507
Encrypted data without ciphertext matching	161736

References at <http://www.sybase.com>

- **White paper: Protecting Personal Data in Sybase ASE**
- **Archived Techcast: ASE Encryption Technology  
Techcast Video**

<http://www.sybase.com/products/informationmanagement/adaptiveserverenterprise/datasecurity>

- **ASE 15.0.1 manual: Using Encrypted Columns in Adaptive Server®**

<http://sybooks.sybase.com/nav/detail.do?docset=899>

- **ASE 12.5.3a manual: New Features Adaptive Server®  
Enterprise 12.5.3a**

<http://sybooks.sybase.com/nav/detail.do?docset=755>

# Agenda

- ✓ Encrypted Columns Feature Description
- **ACS/HCMS Protects Privacy Using ASE**

::ISUG

techcast series

# ACS Human Capital Management Solutions Protects Privacy Using Adaptive Server

Tad Hawkins  
Vice President, IT Architecture, ACS HCMS

SYBASE®

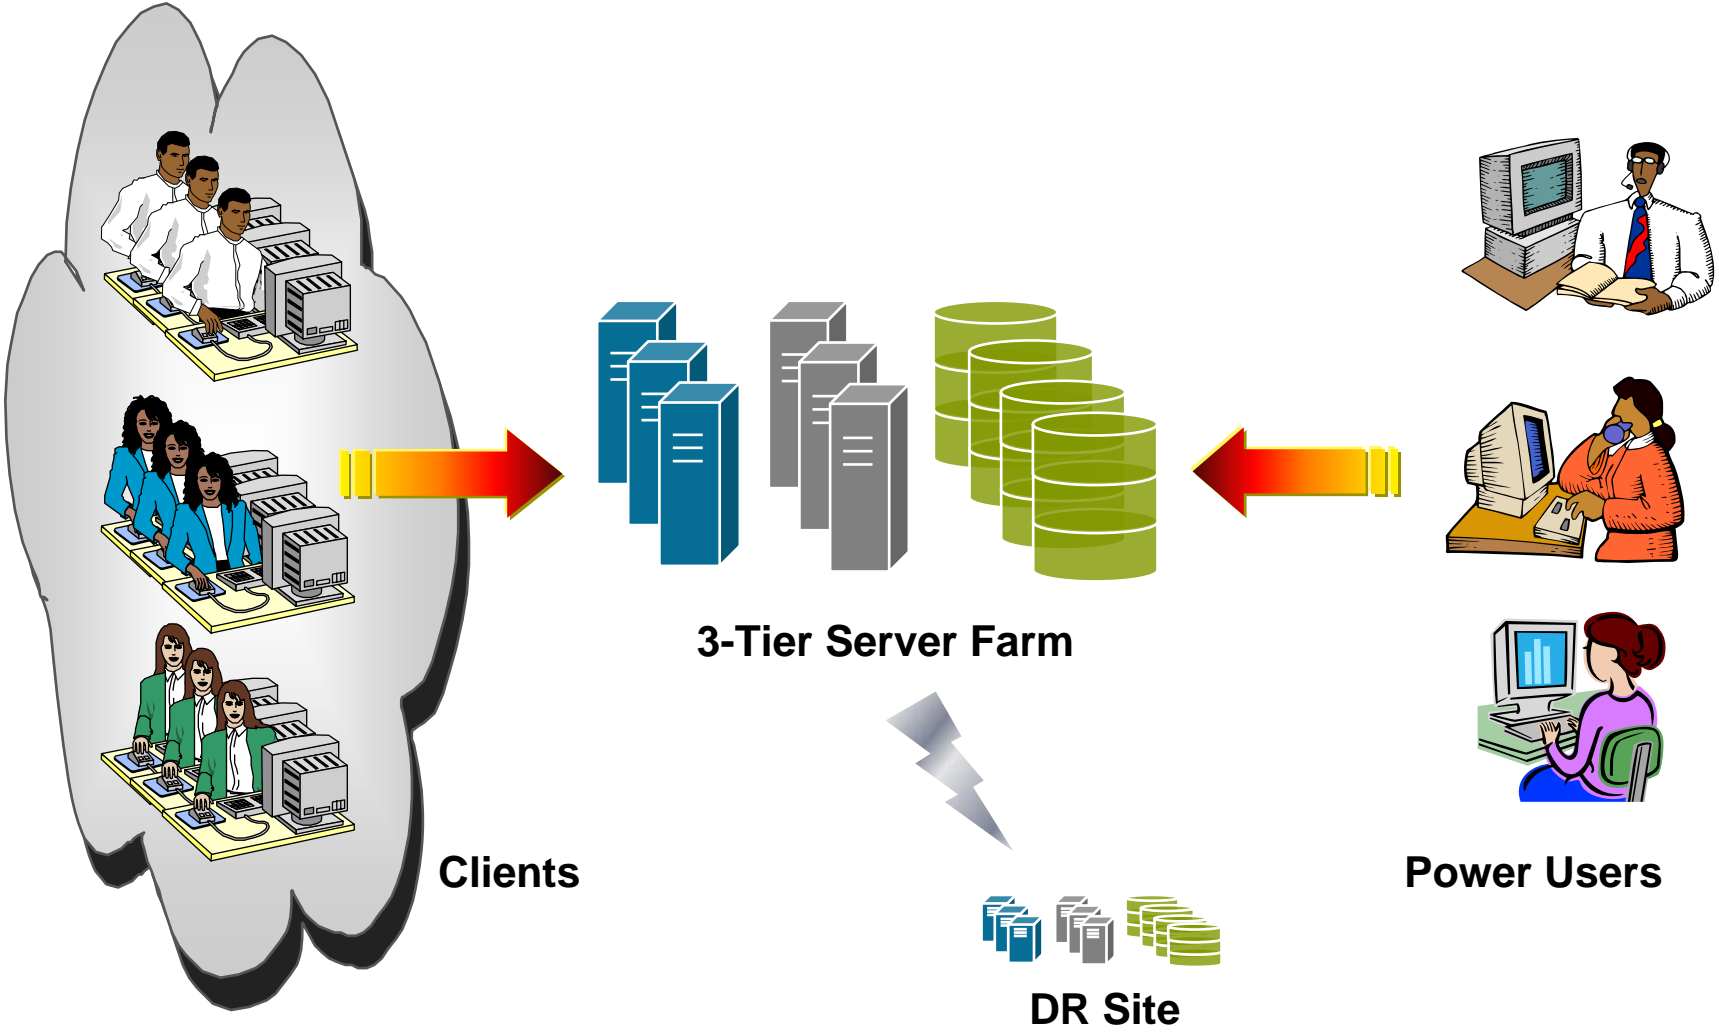
## **Business of ACS: Human Capital Management Solutions**

- **Provides outsourced Human Resources and Employee Benefits administration services**
- **Role: responsible for IT Architecture supporting Employee Benefits outsourcing platform**

## Employee Benefits outsourcing platform

- **Outsourced administration for**
  - 401(k) savings plans
  - traditional pension plans
  - various types of flexible benefit or cafeteria style health and insurance plans
  - deferred compensation plans
- **Platform includes**
  - an employee and manager self-service web site
  - an interactive voice response (IVR) system
  - multiple call centers

# ACS/HCMS: 3-Tier Platform Architecture



## Employee Benefits outsourcing platform contd ...

- **Platform Architecture**

- 3-tier environment: web, application and database server farms
- Supports multiple clients
  - **Client sites share the web and application tiers**
  - **Each database server supports multiple clients**
    - Each client's data resides in client-specific databases
    - RDMS is Sybase's ASE, running on HP UX
    - Replicated in real-time using Sybase Replication Server
- Volume of users
  - **150+ total clients ranging from 3,000 – 250,000 employees**
  - **Approx 5 million potential self-service users**
    - Clients' current and former employees and beneficiaries
  - **Approx 3,000 potential power users**
    - Administrative and call center staff

## Employee Benefits outsourcing platform contd ...

- **Database architecture and size**
  - RDMS is Sybase's ASE, running on HP UX
  - Each client's data resides in client-specific databases
  - Each database contains 273 tables
  - A typical client with 50,000 employees would have
    - **Dozen tables with 1,000,000+ rows**
      - Largest table could be upto 10,000,000+ rows
    - **Few dozen tables with 500,000+ rows**
  - Processes
    - **online real-time transactions**
    - **batch transactions from client's HRIS systems**
  - Peak transactions during enrollment period
    - **Transaction activity: order of magnitude higher than normal**

## Motivation

- **Prior to outsourcing**
  - Clients supported Employee Benefits with in-house systems
  - Did not encrypt data at rest
  - Data includes personal identifying data
- **With outsourcing, data is**
  - Housed outside of the client's main corporate data centers
  - Accessed by an internet facing application
- **Clients' data security standards require encryption of data at rest**

## Requirements

- **Application transparency**
  - Same application used by multiple clients
  - 80% of code is completely generic
  - Changes to generic code to support encryption introduces undesired complexity
- **Performance**
  - Ability to handle 10,000 simultaneous users during last few hours of the enrollment period
  - 182 / 273 tables needed encryption on the identifying column
    - **Identifying column is the leading column for**
    - **Primary keys and of clustered index**
    - **MUST HAVE: Encryption of primary key columns and efficient joins on encrypted columns**
- **Robust key management**
- **Support for disaster recovery**

## First attempt

- **Implemented encryption 3½ years ago with 3rd party product**
  - Met all requirements save one: PERFORMANCE
  - Product supported through-the-year processing levels
    - **After several months of tuning**
  - Could not support peak processing loads
    - **Put encryption in place for most of the year**
    - **Disabled encryption for the fall annual enrollment period**
    - **Continued to look for a solution to keep data encrypted at all times**

## Second attempt

- **Experimented with Sybase's Encryption Option for ASE**
  - Provided feedback on pre and post release versions
  - Initial tests were very encouraging
  - Targeted implementation for fall 2006 annual enrollment season.
    - **Started fairly rigorous and extensive performance testing**
    - **Objective was to convince our client that the solution could handle the peak load**
    - **Worked with Sybase on a couple optimizer issues**
    - **Fixed a few issues within ACS/HCMS code where less-than-optimal coding was only brought to light by the use of the Sybase Encryption Option.**

## Second attempt: Performance test results

Objective	Test Description	Results
Measure encrypted application behavior under peak loads	End-to-end stress tests using recorded web sessions. Simulated 3-times peak load. Replication was enabled.	Same as w/o encryption
Measure encrypted application degradation over time	End-to-end stress tests with sustained peak load activities for 6 hours on each of 3 consecutive evenings. Replication was enabled.	Same as w/o encryption
Measure breaking point for concurrent transactions	Database only tests simulating large number of concurrent users (higher than 3-times peak load).	Same as w/o encryption

- **Met and exceeded SLAs even under 3-times peak load**
- **Surpassed all expectations of ACS**

- **In-production in May and June 2006**
  - 3 weekend production implementation process
    - **Decrypt the prior solution**
    - **Upgrade our ASE Version**
    - **Re-encrypt with Sybase Encryption Option for ASE.**
  - So far so good
    - **In production for four months**
    - **Have run many peak load stress tests during maintenance windows**
  - Looking forward to a smooth annual enrollment process towards the end of October

## **ACS/HCMS: Future Plans**

- **Wait for results from upcoming enrollment period**
- **Move all clients using 3<sup>rd</sup> party encryption to ASE Encryption option**
- **Offer encryption service to other clients**

- **Test. Test again. Test some more!!**
- **Use life-sized encrypted test environment**
  - When encrypting leading index columns
  - When preparing for large OLTP environment
- **Be aware of pitfalls**
  - Use same key to encrypt joining columns
    - **Ran into some issues with cross-database joins**
  - Encrypt columns in temporary tables used in joins with same key as the original tables and columns
    - **May require application changes**
  - Move SARGS using “like” clauses on encrypted columns e.g.  
`AND my_pk_column not like "...%"`  
to end of where clause
    - **Sybase should look into optimizing this automatically**

## **ACS/HCMS: Summary**

**ASE Encryption Option is exactly what ACS needed**

- ✓ **Application transparent encryption facility**
- ✓ **Maintains performance**
- ✓ **Supports disaster recovery environment**

# Agenda

- ✓ **Encrypted Columns Feature Description**
- ✓ **ACS/HCMS Protects Privacy Using ASE**