

Addressing the Challenges of Data Retention Compliance

The Verint/EMC Solution

TECHNICAL OVERVIEW

December 2007



Table of Contents

Preface	1
Solution Overview.....	2
More Complete and Seamless Compliance with Data Retention Directives	2
Superior Scalability	3
Built-In Data Security and Privacy Protection	3
The Solution at a Glance	3
Solution Architecture.....	4
Data Collection and Preparation	4
Data Storage.....	4
Data Model	5
Data Compression.....	5
Data Storage Security	5
High-Availability Storage Solution	5
Backup and Restore	6
Disaster Recovery.....	7
Storage Module Scalability	7
Query and Delivery	8
Administration.....	8
An Illustration of the Solution Architecture.....	8
Typical Hardware Configurations	9

Unauthorized use, duplication, or modification of this document in whole or in part without the written consent of Verint Systems Inc. is strictly prohibited.

By providing this document, Verint Systems Inc. is not making any representations regarding the correctness or completeness of its contents and reserves the right to alter this document at any time without notice.

Features listed in this document are subject to change. Please contact Verint for current product features and specifications.

All marks referenced herein with the ® or TM symbol are registered trademarks or trademarks of Verint Systems Inc. or its subsidiaries. All rights reserved. All other marks are trademarks of their respective owners.

© 2007 Verint Systems Inc. All Rights Reserved Worldwide.

Preface



In March 2006, the European Union Parliament and Council published Directive 2006/24/EC, commonly known as the *European Union Data Retention Directive*. Designed to promote pan-European cooperation in the prevention of terrorism and crime, the directive mandates that communications service providers throughout the EU retain communication-related data — but not communication content — for a period of 6 to 24 months. Such communication-related data includes call records, phone numbers, call duration, locations for cellular calls, sender, recipient, and time of email communications, as well as subscriber information.

This new data retention directive presents communications service providers with some formidable challenges, from securing retained data to complying with the individual data retention regulations of each EU member country. Perhaps most significant, however, is the need to implement an infrastructure that can collect, process, and store hundreds of terabytes of data about the communications that are transmitted over their networks.

To help operators address these challenges, Verint® and EMC have partnered to deliver a powerful, yet cost-effective data retention solution built on Verint STAR-GATE™ DR and the EMC CLARiiON CX UltraScale series storage platform system. This end-to-end data retention solution handles data collection, storage, and administration, as well as data query, retrieval, and delivery to law enforcement agencies in compliance with data retention directives and regulations.

This technical overview explores the benefits of using the Verint/EMC solution and provides an in-depth examination of the solution architecture.

The largest impact to operators is the requirement to store extraordinary amounts of data The data to be stored in compliance with the data retention regulations is expected to be in the tens of terabytes, if not hundreds of terabytes for larger operators, and for much longer terms.

EU Data Retention Directive: An Onerous Burden on Service Providers
Frost & Sullivan, 2006

Solution Overview

The Verint/EMC solution facilitates communications service provider compliance with the European Union Data Retention Directive, specific national data retention regulations, and warrants for communication data. Built on our expertise from hundreds of diverse lawful interception implementations worldwide, this solution is designed to:

- Collect call data records (CDR), subscriber information, and network traffic data from any communication network.
- Retain communication data for the required 6-24 month period.
- Secure data against unauthorized access or use.
- Retrieve and deliver data to the appropriate authorities, as required by law.
- Provide a centralized solution for comprehensive, transparent, multi-country compliance.

The Verint/EMC solution supports centralized, distributed, and hybrid centralized/distributed configurations and is well suited to address global operator requirements.

Supports A Wide Array of Communication Technologies and Services

The Verint/EMC offering includes passive, active, and hybrid solutions for a wide range of communication services and communication technologies.

Wireline

Fiber, xDSL, Cable

Wireless

GSM/GPRS/EDGE/UMTS (3GPP), CDMA2000 (3G), WiMAX (4G), Wi-Fi, Satellite

Services

VoIP (SIP, H.323, MGCP), IMS/NGN, Push-To-Talk (PoC), Instant Messaging, MMS, SMS

More Complete and Seamless Compliance with Data Retention Directives

The Verint/EMC solution helps service providers fully comply with data retention directives without impacting service levels or network operation. The solution uses commercial off-the-shelf hardware and leverages existing IT networks, reducing deployment and maintenance costs and adding value to IT investments.

The Verint/EMC solution can be easily customized to operate with virtually any CDR source or subscriber data repository and can be easily integrated into operator environments. This flexible, feature-rich solution is storage, hardware, and database agnostic — built to work seamlessly with:

- Any communication network
- Leading billing mediation platforms
- Existing IT infrastructures
- Industry-standard hardware, software, and storage solutions

This interoperability streamlines deployment, enhances reliability, and reduces the need for additional capital investments and costly integration services.

Superior Scalability

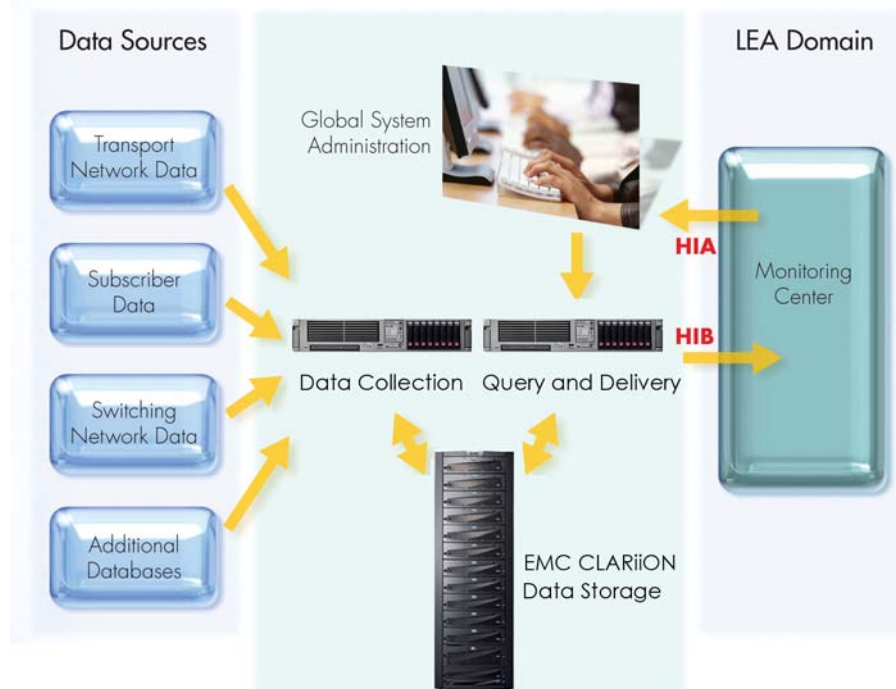
The Verint/EMC solution can manage network topologies from small, single-switch implementations to country-wide deployments. Its modular architecture readily adapts to meet the demands of new communication protocols and technologies and makes upgrades easier. Operators can expand their data retention capabilities or add this joint solution to their existing STAR-GATE Lawful Interception or STAR-GATE Lite solution simply by adding a module to the system already in place.

Built-In Data Security and Privacy Protection

The Verint/EMC solution incorporates a wide array of security features to protect communication data from unauthorized access and help ensure timely and secure delivery to the appropriate agencies. These features include audit trails, flexible and detailed user role definition, data query encryption, transfer communication encryption, multi-tier database access control, operating system and platform hardening, and more.

The Solution at a Glance

The Verint/EMC solution features powerful administration, query, and delivery capabilities, with components that include: high-compression storage functionality for optimizing use of storage resources; multi-platform, multi-technology collection to accommodate a variety of communication network topologies; and high-capacity IP probes for generating communication session flow summaries.



This illustration shows how the Verint/EMC solution interfaces with data sources in the operator's network and the monitoring centers in the law enforcement agency.

The solution interfaces with law enforcement agencies via the ETSI Handover Interface (HIA and HIB), which enables query and standardized delivery of retained call data records and subscriber data information.

Solution Architecture

The Verint/EMC solution contains four modules:

- The **collection and preparation module** interfaces with the network elements and collects the data to be retained in compliance with data retention directives and regulations.
- The **storage module** serves as the data repository, retaining the data for the time period defined in the data retention directives and regulations.
- The **query and delivery module** retrieves the data from storage, as requested by a warrant, and delivers it to the law enforcement agency.
- The **Global System Administrator (GSA)** provides a single point of administration, with an intuitive graphical user interface that enables operators to execute queries, perform system administration, and secure data from unauthorized access or use.

Data Collection and Preparation

The data collection and preparation module collects the data to be retained from the operator's network and pre-processes it as preparation for storage. Functionality includes reduction of fields for storage, conversion of CDRs to unified formats, CDR enrichment, persistency, integrity verification, and more. This module provides support for interfacing with network switches, CRM/ERP databases, and passive probes, receiving data in batches or streamed, and using collection modules from leading billing mediation vendors. The processed data is then sent to the storage module, either through direct marking to a database or as files.

Data is collected from various data sources, such as network elements, probes, CRM systems, information databases, location servers and more. The different data streams are processed to remove unnecessary fields and add fields from other sources, thus creating enriched records. The processed data streams are then forwarded to the Oracle database and stored in the EMC CLARiiON platform.

Data Storage

The data storage module includes an Oracle database using EMC CLARiiON as the storage platform for writing the collected information into the repository and fast retrieval of queried data. The storage module was designed to meet the unique requirements of the data retention system, including:

- Time-based partitioning for efficient retention handling
- Smart indexing for quick retrieval of large volumes of data
- Advanced compression technologies
- Enhanced security
- High availability and scalability
- Backup to remote sites for specific disaster recovery requirements

The storage solution enables data to be stored for diverse retention times, simplifies data backup and deletion of expired records, and facilitates storage scalability by simply adding additional storage volumes. The storage module offers a variety of security-related features for both the stored data (such as data encryption and data integrity assurance) and the storage solution and interfaces (such as WORM/Write-Once-Read-Many, communication link encryption, platform redundancy, and a high-availability architecture).

Data Model

This module is based on a single database instance with dedicated tables in the database scheme for each service, such as voice telephony and emails. Every table is partitioned by time, and a table per day per service is created. The number of partitions reflects the required retention period. The retention period can vary among services (and countries) to meet local regulation requirements. When the retention period expires, the relevant partitions are dropped and cannot be accessed again.

Data Compression

The Verint/EMC solution employs sophisticated compression technologies to optimize use of physical storage and input/output (I/O) resources, providing as much as 70-90% reduction in storage volume.

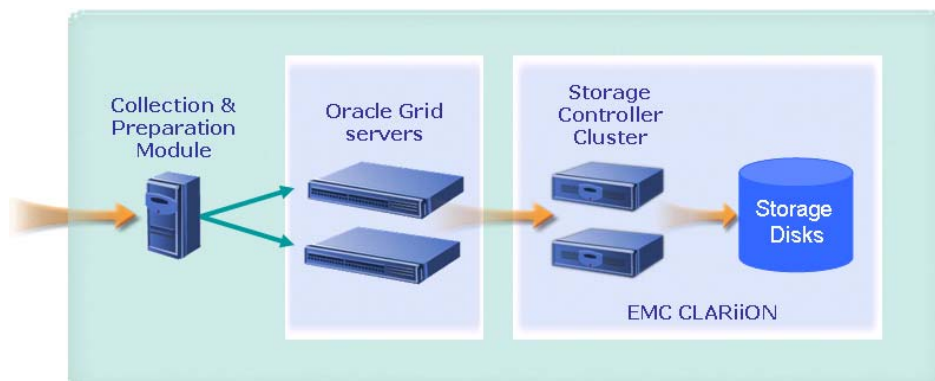
Data Storage Security

The storage module features several security mechanisms, including database authentication, access control to table data at the user and application level, database level audit logs, and communication links encryption.

High-Availability Storage Solution

The Verint/EMC solution may be configured for high availability, with two database servers sharing the same database instance using Oracle 10g Grid capabilities. In addition, EMC CLARiiON achieves five nines of reliability with a combination of state-of-the-art CPUs, high-performing memory subsystem, and dual active independent processors for continuous operations and automated failover.

The storage hardware provides redundancy features, such as duplicate power supplies, dual fans, clustered storage controller, and RAID disk arrays.




This illustration shows the solution in fully redundant configuration in which the database is in a Grid configuration and the storage is in a clustered environment.

Backup and Restore

Backup and restore are based on EMC SnapView software, which leverages the power of EMC CLARiiON networked storage systems to create local point-in-time copies of data that can be used for testing, backup, and recovery operations.

Snapshots and clones speed up tape backup, dramatically reducing the time required to recover data directly from disk drives. EMC SnapView features consistent split capabilities to ensure data coherency for local replication needs with the flexibility to create pointer-based, space-saving snapshots and highly functional, full-volume clones.

EMC SnapView effectively creates a zero-time backup window. This is facilitated by offloading backup cycles from the production systems without disrupting production, giving users the flexibility to conduct backup at virtually any time. Also, snapshot rollback enables snapshots to be restored directly to the production volumes, as illustrated below.



SnapView enhances business continuance and data protection by doing the following:

- Keeping applications online and productive during backups
- Running backups against a logical image
- Allowing more frequent backups for data protection
- Maintaining application integrity through consistency technology
- Facilitating instant restore from disk

Disaster Recovery

The Verint/EMC solution utilizes CLARiiON MirrorView/A, an asynchronous storage-system based product that periodically takes a point-in-time copy of the changes in a logical storage block or logical unit number (LUN) of one storage system and replicates the copy to a LUN in another storage system. MirrorView/A asynchronous replication is best suited for environments that experience a low rate of changes in data, such as in a data retention warehouse. Also, MirrorView/A provides backup without impacting the operational system performance and consumes lower network resources by spreading the transfer times from peak hours to night hours.

MirrorView/A is designed to integrate with Oracle to create consistency clusters, maintaining the integrity of a database that is spread over multiple LUNs. This creates a cost-effective remote replication solution that is ideal for Recovery Point Objectives (RPOs) of 30 minutes or greater. The communication links between main and remote sites can be utilized over Fibre Channel or IP. MirrorView also enables data mirroring capabilities between different generations of CLARiiON storage systems.

Furthermore, CLARiiON offers EMC RecoverPoint/SE, a comprehensive data protection solution that provides operational recovery to any point in time with a Continuous Data Protection license option or Continuous Remote Replication. RecoverPoint/SE protects companies from data loss as the result of such common problems as data corruption, software errors, viruses and end-user errors, while also protecting against catastrophic events that can bring an entire data center to a standstill.

Storage Module Scalability

The required storage volume may increase as the result of growth in the volume of input data and/or regulatory extension of the retention period. The storage module can be expanded either by increasing database capacity based on a scale-out approach using Oracle Grid or by increasing EMC storage volume. The latter can be achieved in two ways:

CLARiiON metaLUN Technology: A LUN or logical unit number is a virtual partition or volumes of a RAID set — in other words a logical storage block made out of a partition over several disks. metaLUN technology enables non-disruptive storage expansion by adding new LUNs and aggregating them into metaLUNs.

CLARiiON Virtual LUN Technology: This technology enables movement of LUNs and metaLUNs within a CLARiiON array without disrupting the application. The CLARiiON CX3 UltraScale series supports tiered storage options within a single midrange array. Exploiting UltraPoint's auto-sense speed-setting capabilities, CLARiiON supports the simultaneous use of 2Gb/s and 4Gb/s drives; for example, offering the 4Gb/s 15,000 rpm high-performance Fibre Channel disk drives and high capacity alongside with the low-price 2Gb/s 7200 rpm SATA disk drives. Virtual LUN technology enables the dynamic movement of data from one tier of storage to a higher or lower tier; for example, from the 2Gb/s Fibre Channel to the 4Gb/s Fibre Channel drives. Data can also be moved from a volume of one RAID type to a volume of a different RAID type to meet the performance requirements of a particular application. Virtual LUN technology can also be leveraged to improve capacity utilization by moving data to underutilized spindles in the array.

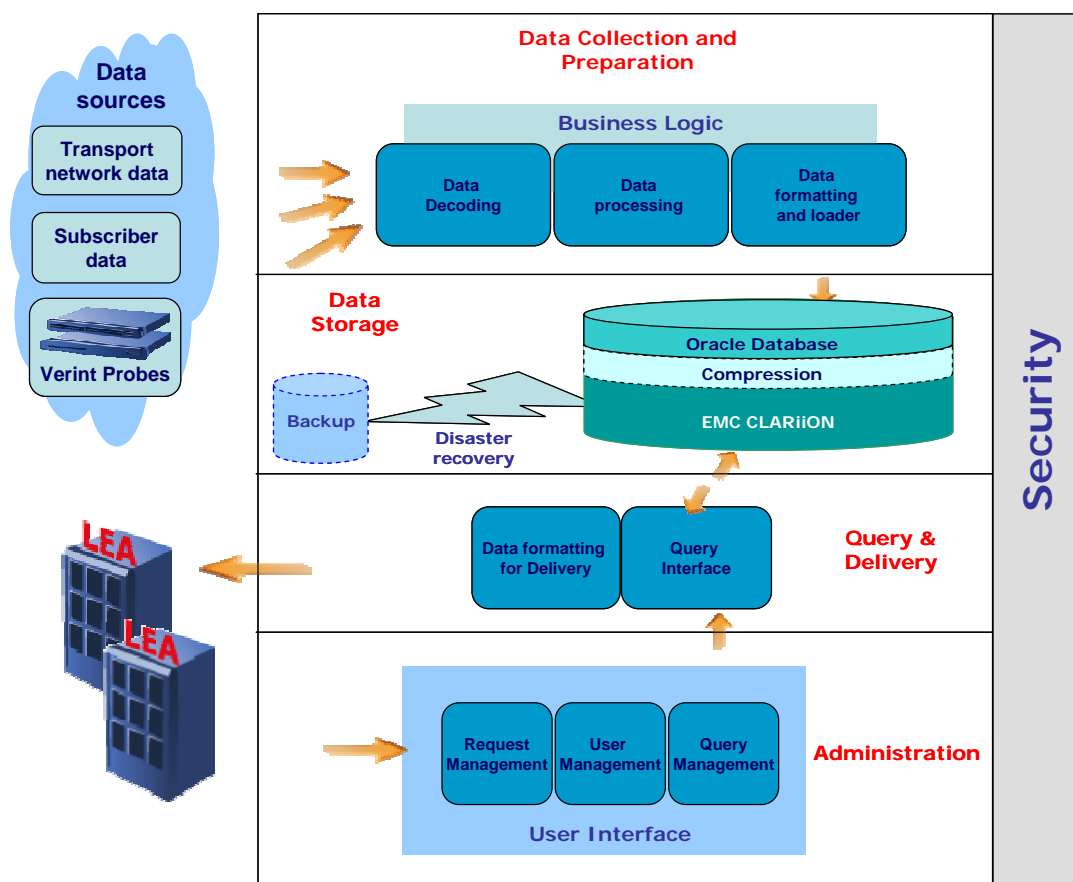
Query and Delivery

The query and delivery module interfaces with the storage module for data retrieval of the required retained data. The retrieved information is then formatted into a CSV or XML format (ETSI delivery format will be supported once it is officially published) and sent to the LEA according to the local regulation or standards. Query response time depends on the volume of data to be retrieved.

Administration

Verint Global System Administration (GSA) features an intuitive graphical user interface for centralized control of user administration, request handling and query management. The GSA is based on .NET Smart Client technology and offers user management with advanced security measures, such as a smart password mechanism and specific user roles and privileges. An audit trail of all activity is created and stored in encrypted mode. The GSA can also serve as a unified administration platform for data retention and lawful interception solutions.

An Illustration of the Solution Architecture



Typical Hardware Configurations

The following table shows the recommended EMC hardware for Verint/EMC solutions for networks with up to 10M/100M/300M/1000M CDRs per day and a retention period of 12 months.

Item	Up to 10 Million CDRs/day	Up to 100 Million CDRs/day	Up to 300 Million CDRs/day	Up to 1000 Million CDRs/day
Storage Platform	EMC CLARiiON CX3-10 1 x Shelf of 146GB HD 1 TB storage	EMC CLARiiON CX3-20 7 x Shelf of 146GB HD 10 TB storage	EMC CLARiiON CX3-40 10 x Shelf of 300GB HD 30 TB storage	EMC CLARiiON CX3-80 32 x Shelf of 300GB HD 100 TB storage

Verint. Powering Actionable Intelligence.®

Verint Systems Inc. is a leading provider of actionable intelligence™ solutions for an optimized enterprise and a safer world. More than 5000 organizations in over 100 countries rely on Verint solutions to perform more effectively, build competitive advantage, and enhance the security of people, facilities, and infrastructure.

About EMC

EMC Corporation (NYSE: EMC) is the world's leading developer and provider of information infrastructure technology and solutions that enable organizations of all sizes to transform the way they compete and create value from their information. Information about EMC's products and services can be found at www.EMC.com.