



PO BOX 984, BOWLING GREEN, NEW YORK, NY 10274. 1-877-659-4309

**Data
Management
Association**

**Visit us at:
www.dama-ny.com**

News Letter

March 2007

DAMA NY OFFICERS

- Denis Kosar**
President
- Beth Ouellette**
VP, Marketing
- Sam Patel**
Program Chair
- Stu Rosenthal**
Membership Chair
- Patricia Komar**
Treasurer, VP Finance
- Chhaya Kapadia**
Chapter Secretary
- Ho-Chun Ho**
News Letter Editor
- Jae Cha**
VP On-Line Services
- Toshu Kabessu**
VP Special Events



Mark your calendars for: **Thursday, March 15, 2007**

**Challenges to make an EDW environment
SOX compliant**

- Awais Inam**
- Imran Zafar**
- Techlogix**

**Location: Computer Associates
520 Madison Avenue, 22nd floor, New
York City, New York**

*Please note that the location has changed. It
is between 53rd and 54th street.*

CONTENTS

Next Meeting- March 15	1
Presentation Abstract	2
About the Speakers	2
Feature Article	3 - 4
Member Discounts	5
Patrons of DAMA NY	6 - 7

AGENDA:

8:30 am - 9:00 am	Registration
9:00 am - 10:00 am	Awais Inam Imran Zafar
10:00 am - 10:15 am	Coffee & Refreshments
10:15 am - 11:45 am	Awais Inam Imran Zafar
8:30 am - 9:00 am	Registration

**To register just go to our website at
[http://www.dama-ny.com!](http://www.dama-ny.com)**

Presentation Abstract

This presentation will provide the audience with an understanding of what Sarbanes Oxley is, as well as the different sections that make up the regulatory requirements. Also we will identify some of the challenges in trying to implement SOX while providing support of an Enterprise data warehouse. We will also identify how to make an Enterprise Data Warehouse SOX compliant and how to prepare for a SOX Audit. The presentation will consist of the following sub topics:

- What is SOX Features and its Organization
- Features of EDW that contrasts with SOX compliance
- How to make an EDW environment SOX compliant
- Preparing for SOX Audit

About the Speakers

Awais Inam, Lead Presenter

Awais Inam is a Senior Consulting Engineer with Techlogix and has worked as an architect and project manager with many different technologies including Oracle, Business Objects, Java, and Teradata. In his most recent role he has been leading a SOX compliance engagement with a Fortune 50 customer of Techlogix.

Imran Zafar, Supporting Presenter

Imran Zafar has more than a decade of software engineering experience and is head of the Business Intelligence/Data Warehousing practice in Techlogix. He primarily works as a BI architect with a special focus on architecting executive reporting/dash boarding solutions.

About Techlogix:

Techlogix is a ten year old software engineering services firm specializing in enterprise solutions in the areas of Business Intelligence/Data Warehousing, Business Process Management, Master Data Management, Integrations, and IT Governance.

Feature Article

17 Simple TIPS for Enhancing Data Warehouse Performance

By Syscore Solutions International 

Most organizations are challenged with the ever-increasing volume and depth of historical data and a need to consolidate it and report on it in real time. This is especially visible in the financial arena because of forces such as compliance regulations, deadlines and introduction of newer, more innovative product lines. The quest for real time data warehousing means reducing the latency between the time when the source data is created, and when it is available for analysis. The following tips can help in building or achieving high performance data warehouses:

1. As with any tuning activity first establish goals and then a roadmap to achieve those goals, improving performance can be a never ending activity, without a starting and endpoint.
2. Strive to localize data that is commonly used together. This will allow your queries to just hit one partition for all joins and removing any need to go across partitions In Oracle terms strive for partition elimination and partition wise joins.
3. Religiously manage database partitions. For example, an Oracle database performs well when the number of partitions and sub partitions is less than 3000 and each partition size is less than 2GB. An overzealous partitioning strategy can bloat the system tablespace.
4. If there are memory constraints, pre-populate small dimension tables in each partition to avoid cross partition joins.
5. Use bitmap indexing for dimension tables, and ensure star transformation of joins to fact tables.
6. Where appropriate, referential integrity should be managed by application logic. This allows loads to perform faster.
7. Keep the number of indexes to bare minimum on large volatile tables, this improves the loads and decreases index re-orgs.
8. Where possible, enforce quality in the database through data uniqueness and domain constraints. Create indexes post load without transaction logging.
9. Take advantage of bulk data facilities in your design wherever possible. Disk Mirroring and transportable tablespaces can be utilized to move large amounts of data efficiently.
10. Be mindful of networking alternatives and that local file access is preferred to access over a network. When shared access is desirable, ensure the best transport available is used.

11. Sometimes straight forward transformations may be done faster outside the database in a program environment that doesn't have the overhead of multi-user concurrency.
12. Use smart surrogate keys for common dimensions like date. For example, using the INTEGER format of the date value as a surrogate key will remove the need to join the fact and date dimension tables to get the date surrogate keys, and this saves a lot of processing power. You can still assign a dummy record with a special surrogate key (such as 0) for all invalid date fields in the fact table.
13. Evaluate explain plans for slower/larger queries to ensure there are no full table scans and lowest possible physical reads. If enough use cases call for aggregation create summary tables...
14. Do not allow dynamic queries, if possible, monitor and capture long executing queries for optimization. Resource limits on end user accounts should govern resources consumed by ad-hoc queries.
15. Model the physical database in a way that is conducive to receiving data in increments throughout the day. Intra-day loading makes data available sooner. Where possible, try reaching directly to the source systems or taping into a message bus to receive source data in real time.
16. A key factor of the data warehouse will be its ability to scale gracefully. Scheduled batch loads and reports should use grid resource virtualization to partition the workload.
17. If full tablescans are to be expected, use a larger block size that is appropriate for the OS block-size to expedite I/O operations.

About SysCore

SysCore focuses on information Integration. Our deep business knowledge coupled with our IT thought-leadership, architectures, frameworks and practical experience in Global Information Architectures, Data Integration, Business Intelligence, and Data Warehouses, provides our clients with the competitive advantage and differentiator that businesses look for when partnering with a services firm. For more information, contact us at info@syscore.com or 212.619.5200

DAMA NY Member Services and Discounts

PMP Certification Program

The Project Management Professional (PMP) Certificate Program is a PMBOK approach to Project Management Certification (35 hour program over 5 weeks). Though the number one objective is to obtain the PMP certification, the participants will gain knowledge, tools, tricks and processes to improve the day-to-day management of tasks, projects, programs, and even organizations.

Contact training@OuelletteGroup.com for more information.

Detailed information is available on <http://www.OuelletteGroup.com>.

Morgan Kaufmann Publishers

20% discounts in books in their data management line. Please use discount code 50MKP when ordering.

Please visit <http://www.mkp.com/datamanagement>.

Meta-Data Professional Organization

The Metadata Professional Organization (MPO) is a newly formed non profit group dedicated to providing a forum where organizations and individuals can find assistance in addressing their metadata issues and strategies, be advised and informed of developments in the field, and share experiences with other practitioners.

For further information please write membership@metadataprofessional.org.

DebTech International and Wilshire Conferences

DebTech International, L.L.C. is a training and development company specializing in corporate onsite and public IT and data management education. DAMA members may attend DebTech events at a discounted rate.

For additional information please email davida@debtechint.com or call (973) 379-7212.

PATRONS of DAMA-NY

We wish to thank the organizations listed below for their support!



Computer Associates

Acclaimed business software portfolio for:
enterprise management, security, storage, business transformation and integration, portal and knowledge management, and predictive analysis

www.ca.com

Michelle Accardi

Michelle.Accardi@ca.com

Tel: 631-342-4236



Syscore Solutions International

The Leader in providing Information Management, Data Warehouse, Middleware Architecture and Business Solutions

Bruce Gallager, Fax 212-292-5093
Email: BruceG@Syscore.com



InfoModel, Inc.

The Leader in Database and Data Warehouse Training and Solutions

Tom Haughey
201 337 9094, email:
thaug@attglobal.net



Siperian
Uniting Customer Data with Business Processes.

Gregg Loos
Director of Sales
Tel: (917) 533-3640
GLoos@Siperian.com

DebTech International L.L.C.

Education By Industry Leaders
Seminars and Conferences for Information Management Professionals

<http://www.debtechint.com>
David Berger (973) 379-7212
info@debtechint.com



GoldenSource

*GoldenSource is a proven software solution based on the industry's most comprehensive global financial data model incorporating over \$150M in research and development and the company's 20 years of **experience.***

Joseph Cubias
Marketing Manager
Tel: 212.798.7174
jcubias@thegoldensource.com



ChoiceMaker Technologies *When Accuracy Counts*

ChoiceMaker is a data quality company which uses patented, machine learning techniques to build ultra accurate systems for matching records and removing duplicate records from databases.

Robert Fragola
Vice President Sales and Marketing
Tel: 212.918.4411
robert@choicemaker.com



Odyssey

Solutions for private banks, asset managers, wealth managers and large brokerage firms worldwide Odyssey Asset Management Systems, Inc. is one of the leading providers of private banking, wealth & asset management technology worldwide. Over 200 financial institutions in the U.S., Canada and 27 other countries use Odyssey applications for SRM (Securities Reference Master), Portfolio Management and Client Relationship Management.

Claire Frankel
Senior Pre-Sales
Consultant
+1 917 699 9024
cfrankel@odyssey-
group.com

Candyce Cohen
Senior Sales Executive
+1 201 600 1083
ccohen@odyssey-
group.com



IBM

Information on Demand

The IBM WebSphere® Information Integration Solutions platform integrates and transforms any data and content to deliver information you can trust for your critical business initiatives. It provides breakthrough productivity, flexibility and performance, so you and your customers and partners have the right information for running and growing your businesses.

Bill Roberts
District Sales Manager
Information
Integration Solutions
IBM Software Group
212-745-3337
roberts2@us.ibm.com

Ellie Stone
WW Events Team
IBM Information
Management
IBM Software Group
508-366-4858
elliest@us.ibm.com

Become a Patron

Your company's information can be posted here!

For information about the Patron program please call (877) 659-4309

For information about the Patron program call Denis Kosar 1-877-659-4309