

Sponsored by:

SUNGARD®


SmartStream



Volanté 
The Messaging
Company

With more than 1,500 clients, including the world's top 10 banks*, we know what it takes to support the industry's leaders.



Redefining post trade operations globally

SmartStream delivers highly flexible and innovative transaction processing software that supports critical operational infrastructures: from industry-leading reconciliations and cash management solutions to reference data management services.

Our TLM® solutions enable you to rapidly respond to changing market demands – where efficiency, risk, cost and compliance are the new industry watch words.

See what we can do for your organisation, contact SmartStream today:

+1 (866) 603-1893 (Americas)

+44 (0)800 279 5465 (UK)

+65 6 224 76 89 (Asia Pacific)

info@smartstream-stp.com

Card Processing
Cash & Liquidity Management
Compliance Management
Corporate Actions Processing
ETD/OTC Derivatives Processing
Fees Invoice Management
Reconciliations
Reference Data Management
Trade Finance
Trade Process Management
Treasury Confirmations


SmartStream

smartstream-stp.com

* The Banker, July 2010

Editor-in-Chief Victor Anderson
victor.anderson@incisivemedia.com
tel: +44 (0) 20 7316 9090

Online and US Editor Rob Daly
rob.daly@incisivemedia.com
tel: +1 646 490 3968

Deputy Editor, Buy Side Anthony Malakian
anthony.malakian@incisivemedia.com

Deputy Editor, Sell Side Michael Shashoua
michael.shashoua@incisivemedia.com

European Reporter Sitanta Ni Mathghamhna
sitanta@incisivemedia.com

Head of Editorial Operations Elina Patler
elina.patler@incisivemedia.com

Contributors

Max Bowie, Editor, Inside Market Data
Tine Thoresen, Editor, Inside Reference Data

Publishing Director Lee Hartt
lee.hartt@incisivemedia.com
tel: +44 (0) 20 7316 9443

Commercial Director Jo Garvey
jo.garvey@incisivemedia.com
tel: +1 646 736 1840

US Business Development Manager Melissa Jao
melissa.jao@incisivemedia.com

European Business Development Manager
Robert Alexander
robert.alexander@incisivemedia.com

Senior Marketing Manager Claire Light
claire.light@incisivemedia.com

Design Lisa Ling
lisa.ling@incisivemedia.com

Chief Executive Tim Weller
Managing Director John Barnes

Head Office
32-34 Broadwick Street
London
W1A 2HG, UK

Incisive Media US
55 Broad Street, 22nd Floor, New York,
NY 10004- 2501
Tel: +1 646 736 1888
Fax: +1 646 390 6612

Incisive Media Asia
20th Floor, Admiralty Center, Tower 2, 18 Harcourt Road,
Admiralty, Hong Kong, SAR China
tel: +852 3411 4888, fax: +852 3411 4811

Subscription Sales
Gillian Harker Tel: +44 (0)20 7968 4618
Dominic Clifton Tel: +44 (0)20 7968 4634
waters.subscriptions@incisivemedia.com

Incisive Media Customer Services
c/o CDS Global, Tower House
Sovereign Park, Market Harborough
LE16 9EF, UK
Tel: 0870 787 6822 (UK)
Tel: +44 (0) 1858 438 421 (overseas)
incisivehv@subscription.co.uk



Expanding the data toolbox

There is an old saying that goes: "If you only have a hammer, everything looks like a nail." Basically, if you tie yourself to one approach, your results are going to be very limited at best.

In today's market, financial services firms are finding themselves having more in common with search engines, social media sites and national governments: they all need to manage and analyze exabytes and petabytes of data on a regular basis, but the trusty relational databases that they have relied on for so long are just not keeping up with their users' demands to manage "big data" these days.


Twin forces are driving the desire for improved data management and the enterprise scale. First, there are new ways to analyze structured and unstructured data that previously had not been collected in order to find new potential alpha, such as social media sentiment and other unstructured data. And second, there is the increasing demand from regulators that requires market participants provide more transparency around their transactions by providing more data to the regulators.

In this special report, *Waters* speaks with leading broker-dealers and vendors to discuss how firms can address these issues and what they would like to see in the next generation of data management tools, and touches on the best strategies to implement data integration and standardization. ■

Rob Daly, Online and US Editor

Waters (ISSN 1068-5863) is published monthly (12 times a year) by Incisive Media. Printed in the UK by Wyndeham Grange, Southwick, West Sussex.

©Incisive Media Investments Limited, 2011. All rights reserved. No part of this publication may be reproduced, stored in or introduced into any retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written permission of the copyright owners.



Inside Market Data Inside Reference Data



waterstechnology

Cisco Optimizes Switches for Low-Latency Data Fabrics

Networking and technology vendor Cisco has unveiled optimized switching hardware designed for low-latency market data delivery that forms a key component of its new High-Performance Trading Fabric—a fully optimized infrastructure for financial data and trading communication.

Paul Jameson, global director of financial services at Cisco, says the vendor's new 3064 and 5500 switches, which Cisco began shipping to clients in early May 2011, have been designed specifically to support ultra-low latency and peak message volumes, as a critical component of the trading lifecycle.

"We built features into the switches to optimize them for high-performance trading in financial markets, including features for managing and optimizing data across Layer 2 and Layer 3, allowing institutions to optimize various layers of their high-performance environment," Jameson says. "Without an integrated approach, optimizing performance and upgrading individual areas is very difficult."

Activ Speeds ActivFeed

Activ Financial has released version 1.5 of its ActivFeed low-latency consolidated data feed. The update features reduced bandwidth utilization via a more bandwidth-efficient protocol for updating data; improvements around implementing updates and changes to feed handlers; more comprehensive record information that explains whether an order was filled or cancelled; datacenter redundancy; and lower latency through the use of variable size datagrams, as well as simultaneous use of multiple compression algorithms, and explicit time-based flush.

LSE Licenses Last Sale Prices to Websites

The London Stock Exchange (LSE) has launched a real-time streaming trade data service aimed at websites, which will allow licensed redistributors to broadcast last sale prices for London equities online.

US-based global business news portal CNBC.com will be the first to take up the service, and will begin disseminating last sale prices free of charge to coincide with the launch of a European version of

its news portal. The license will allow CNBC to publish a subset of the LSE's Level 1 real-time data feed without charging end-user data fees, and without requiring users to enter an official LSE login or user session.

Last trade content will be limited to instrument name, instrument identifier—such as a stock's ISIN and/or SEDOL number—last traded price for that

instrument, the volume of shares executed in that transaction, and the time of the trade. While the Last Trade Price data does not have the same breadth as the LSE's unbundled post-trade package, which it released in October last year, officials are confident that the service will provide value to licensees such as CNBC, which can package it into their broader services with news, data and commentary.

Exegy Expands Feed Handler Line-up

St. Louis, MO-based hardware ticker plant vendor Exegy is developing a range of new feed handlers in response to fragmentation in Canada and Asia-Pacific driving increased interest among high-frequency trading clients.

Exegy has begun developing a feed handler for full-depth market data from TMX Select—a new alternative trading system operated by TMX Group that is scheduled for launch this month—and plans to make the TMX Select feed handler available from the first day of

the venue's launch to enable clients to begin processing the data immediately, says Jeff Wells, vice president of product marketing at Exegy.

Exegy is also developing a feed handler to process the Australian Securities Exchange's MarketPoint feed of equities and futures price data, as the pending launch of Chi-X Australia later this year marks the start of fragmentation in Australia, generating more interest from Exegy's high-frequency trading clients.

Interactive Data Enhances PlusFeed

Interactive Data has announced that it has expanded its coverage of Middle East feeds via PlusFeed, the date specialist's low-latency consolidated data feed. Real-time pricing data from the following exchanges is now live via PlusFeed: Abu Dhabi Securities Exchange, Bahrain Stock Exchange, Dubai Financial Market, Kuwait Stock Exchange, Muscat Securities Market, Nasdaq Dubai, Qatar Exchange, and Saudi Stock Exchange.

The vendor's PlusFeed is a high-capacity, low-latency data feed that delivers global, multi-asset class instrument coverage from more than 450 sources worldwide, including over 125 exchanges, with extensive Level 2 data. The feed is used by financial institutions, hedge funds and execution management system (EMS) providers worldwide to power algorithmic and electronic trading applications, as well as many other time-sensitive applications.

waterstechnology

Premium pickings

Do you want to cherry-pick the most valuable content and tools for your company in minimal time?

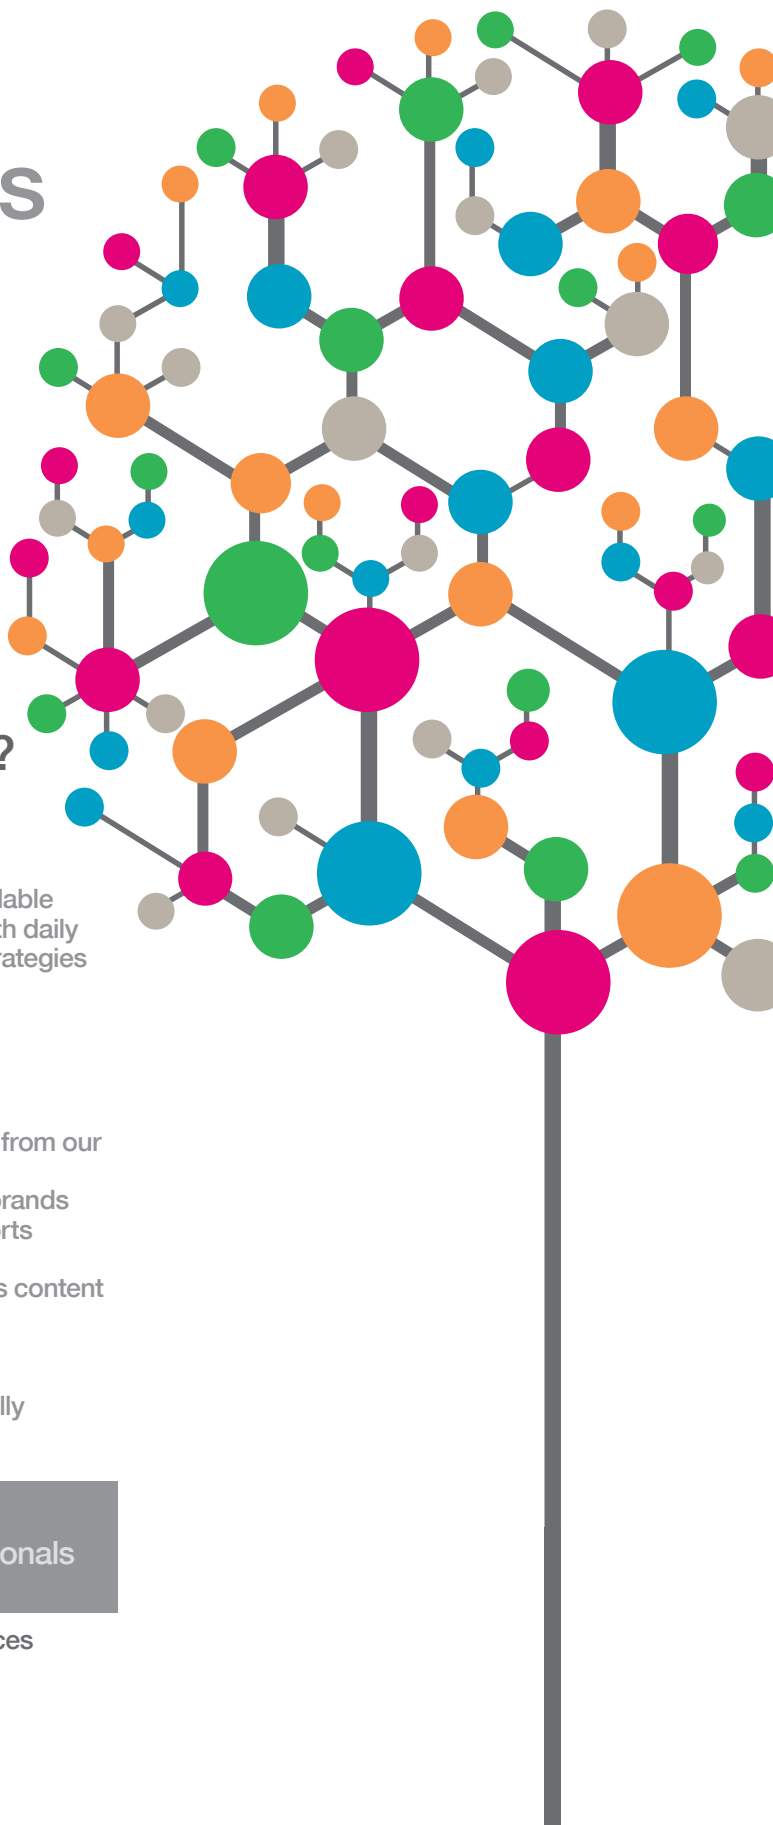
Our brands have long been the choice for professionals within financial-market technologies. Our premium content from Waters, Inside Market Data, Inside Reference Data, Buy-Side Technology and Sell-Side Technology is now available to you via a multi-channel business intelligence platform with daily analysis and news to enable businesses to deliver better strategies with more efficiency than before.

By integrating our five market-leading brands you will get:

- the industry's most respected editorial teams under one web interface
- up-to-the-minute analysis and news – uploaded directly from our editorial desks
- easy navigation – articles linked by topic across all five brands
- free content from our extensive selection of special reports and webinars
- comprehensive news alerts delivering premium business content
- consolidated events and training calendar
- one community to network and interact with
- a multi-level premium subscription service to serve your company's needs – individually, departmentally or globally

waterstechnology.com
Premium content for financial-market professionals
Trial today – visit waterstechnology.com/trial

For details about individual access and corporate site licences to the entire WatersTechnology platform, contact waters.subscriptions@incisivemedia.com



Burststream Preps ‘Fast Normalized Data’

New York-based low-latency data start-up Burststream is preparing to launch a managed service that will provide ultra-low latency access to market data from US equities and derivatives exchanges in the second half of this year.

Burststream, which had the working title Fast Normalized Data, will use instances of French data technology vendor NovaSparks’ hardware-based low-latency feed handlers and ticker plant, co-located in five key data centers that house matching engines or

gateways to the major US markets—Verizon’s facility at 1400 Federal Boulevard in Carteret, NJ, which houses Nasdaq’s matching engine; 165 Halsey Street in Newark, NJ; Equinix’s NY4 facility in Secaucus, NJ; Savvis’ NJ2 data center in Weehawken, NJ; and Telx’s data center in the 350 East Cermak Road facility in Chicago. Clients will be able to either take rack space in Burststream’s cages or cross-connect to the vendor’s NovaSparks appliances from their existing co-location setups.

Desjardins Looks to Eagle for IFRS Support

Montreal-based Desjardins Asset Management, a subsidiary of the Desjardins Group, has tapped Eagle Investment Systems for its data management platform, Eagle Access. In light of new financial reporting standards in Canada, Desjardins signed with Eagle because the vendor had introduced new International Financial Reporting Standards (IFRS) support back in 2009, and had an established product. Desjardins helped with that build-out.

“Eagle’s innovative technology and continued investment in research and

development are reasons we selected Eagle,” says Gregory Chrispin, chief operating officer at Desjardins Asset Management. “A significant aspect of this decision was based on our collaboration with Eagle to support its build-out of IFRS. Our previous technology was not current with IFRS or the instrument coverage necessary to grow our business, so it was important for Desjardins Asset Management to move onto a new data management platform and investment accounting solution.”

Vegasoul Deploys QuantHouse

French low-latency data vendor QuantHouse has announced that Hong Kong-based hedge fund Vegasoul Capital Management is to roll out its QuantFeed low-latency consolidated data feed solution and its QuantLink proprietary network to support the firm’s trading globally.

Trade Associations Release Timeline for LEI Process

The coalition of financial services trade associations reviewing legal entity identification requirements have officially launched a formal process to review solution providers, including standards and registration authorities. The group is working to a tight deadline, and asked solution providers to register their intent to respond by May 18. The complete response is due on June 3, and providers will then be invited to present their solutions by June 17.

According to the current time schedule available on the Sifma website, the final recommendation will be published on July 8.

Telekurs, Xtrakter Ally on Pricing Front

SIX Telekurs has begun distributing pricing data for more than 45,000 fixed-income instruments, including illiquid assets such as asset-backed securities, sourced from trade matching and reporting provider Xtrakter’s XM2M mark-to-market valuation service and bid/offer, high/low/mid-point traded prices from its Trax trade matching service.

Demand from Asia for Access to Feed Handlers

ASX is the first Asia-Pacific venue for which Exegy is developing a feed handler, and Jeff Wells, vice president of product marketing, says the vendor also expects to see demand from clients for data from the Tokyo Stock Exchange, following last year’s launch of its low-latency Arrowhead trading platform.

Exegy has also moved its feed handlers for OTC Markets’ multicast feed and BATS Global Markets’ US Equity Options feed into beta testing. BATS Options data was already distributed as part of the Options Pricing Reporting Authority feed, but as the market has grown, so has client demand for direct access to the

data, Wells says. The vendor is also developing feed handlers for NYSE Euronext’s Global Index Feed and Credit Suisse’s Light Pool ECN, and data from the Mexican Derivatives Exchange—a result of MexDer’s order-routing partnership with CME Group.

“As the market has grown, so has client demand for direct access to the data.”

Data Management: The Building Blocks of Regulatory Reporting

Recent regulatory developments are the single largest event taking place in our industry today. As a result of the economic crisis, regulators around the world are making sweeping changes, not just in relation to reporting but around how financial institutions will be expected to do business in the future. It's not surprising that every industry working group is responding to numerous requests for comments from around the world. **By**

Tony Scianna



Tony Scianna
SunGard

Market participants are tracking a broad range of developments, including the shift of over-the-counter derivatives to a cleared model; collateral management optimization; new trade reporting requirements; and the creation of new bodies such as Swap Execution Facilities (SEFs) and swap data repositories (SDRs). Financial services firms need to prepare—from both rules that have been announced and those that are still being devised.

The financial markets are still analyzing the scant details of the new financial and accounting rules, which include Dodd-Frank, Basel III, International Financial Reporting Standards, and the European Securities And Market Authority (ESMA, formally the Committee of European Securities). However, both new and existing reporting organizations, including the newly formed European Banking Authority (EBA), ESMA, Office of Financial Research (OFR), and the SDRs will require new and more frequent reporting.

A strong enterprise data management framework will be fundamental to being able to manage with whatever develops. This framework will require firms to:

- Capture information in as close to real time as possible.
- Consolidate data across disparate applications.
- Have unfettered real-time access to all of the information.

Capture the Data

Many financial services firms rely on a series of legacy applications that are based on geo-

graphical location and/or asset class. Data is stored in multiple systems, formats, architectures and technologies. Often the environments are batch-oriented or involve end-of-day processes that do not support the output of relevant information until after the processing is performed. Moreover, for global companies there is always at least one system in batch process.

Given the new regulatory and business environment, those tasked with determining a true enterprise data management structure have realized that this is not an acceptable methodology. Somehow data must be captured from multiple source systems across asset classes, and geographical locations on an intra-day basis in addition to any existing nightly and batch processes.

Consolidate the Data

Once this disparate data has been captured from siloed systems, the question becomes, How do you consolidate data that is in multiple formats, technologies and architectures? It is of no use in its current state because it cannot be utilized holistically. To make it usable, the data must be standardized, normalized and stored in a single repository. The repository must be able to support the creation of real-time holistic views as required by senior management or regulators.

This approach is critical to successfully responding to, and even anticipating, market events and new regulations.

Access the Data

After it has been captured, standardized and normalized, the data must be readily accessible in real time. The repository needs to be available 24 hours a day, allow users to run ad hoc queries and be able to respond in any format required.

Firms can then utilize this data to formulate responses to reporting requirements and feed any other downstream application that requires holistic information. This will also help firms answer questions from senior management in a timely fashion.

Financial services organizations will continue to face a series of changing requirements for years to come. And while it may not be the subject of every headline, employing an effective data management strategy is one of the most critical factors in successfully navigating the post-reform landscape. A flexible but comprehensive approach to data management to capture and consolidate data and make it accessible will help firms be ready for both today's challenges and the next wave of regulations, rules and business opportunities. ■

Tony Scianna is deputy head of strategy and marketing for SunGard's Position, Risk and Operations business. Follow him on Twitter @tonyscianna or visit www.sungard.com/streamforposttrade.



Managing the Data Deluge

As data warehouses grow from being measured in terabytes to petabytes and beyond, their owners are faced with the challenges of managing structured and unstructured data on a scale inconceivable just a few short years ago. Those who can manage to analyze these massive quantities of data and act on the results are expected to reap untold rewards, but first they have to find the right tools for the task. *Waters* speaks to seven data management insiders about the state of the tools currently available to the data management industry

Q Are the tools currently available for data management keeping pace with the growth of data that they are managing? How do they need to change?

Rupert Brown, principal architect, Bank of America Merrill Lynch:

No. The large IT giants like IBM, Oracle etc generally behave myopically to defend the suite of data capabilities, such as database management systems (DBMSes), middleware and other platforms that they have acquired over time. Somehow we need to re-segment the industry to create transversal data specialists, independent of the database and middleware vendors.

Michael Tobin, managing director, Knight: On the software side, vendors continue to make incremental improvements to the tools they offer. On the hardware side, where the speed of storage—ie hard disk drives—has not kept pace with increasing volumes of data, the advent of flash memory-based storage systems over the past couple of years has given substantial increases in speed. We have seen in some cases up to a 90 percent reduction in times to produce reports on flash memory-based systems. This performance improvement has permitted us to provide more interactive tools to users, which were historically too slow when accessing large volumes of data on standard disk arrays.



Michael Tobin
Knight

John Jacobs, COO and director of operations, Lime

Brokerage: In general I would say yes, but with certain exceptions. There are more and more tools becoming available, but they are centered around delivering data to end-customers and not as focused on facilitating data delivery by intermediaries and redistributors. For example, NYSE Euronext has been advertising its SuperFeed, and there is a burgeoning plethora of field programmable gate array (FPGA) feed handler providers. However, there seems to be a lack of third-party tools centered on helping to deliver the service for intermediaries, especially toolkits that are integrated into the high-performance solutions for market data user and entitlement tracking.

Robert Brachowski, reference data product manager,

Eagle Investment Systems: My answer is a qualified yes. A successful data management practice recognizes data management as an iterative process, requiring continual improvement and updates for future changes. The same can be said for data

management solutions—the best ones provide continual improvement driven by deep market knowledge and significant R&D budgets, or they can quickly become irrelevant.

Due to the increased regulatory pressure on companies both in Europe and the US, we have seen a tighter integration between data management and business process. This, along with the continued economic environment, is causing clients to ask their data management solutions to do more than they have in the past. Data management solutions must remain flexible in order to solve the businesses' constantly evolving needs. Data management tools must now be focused on the entire lifecycle of data: integrating data from multiple sources, enriching data and ensuring the quality of the data, and delivering the data to the correct consumers. The data management tools must be able to provide these functions for all types of data, including reference data, portfolio-level data, and enterprise data.

Vijay Oddiraju, CEO, Volante: When it comes to handling massive throughput with dependable speed, the efficiency of the fundamental data architecture becomes the crucial issue. This is about absorbing the right inbound data into consuming systems without glitches, expediting downstream flows, and optimizing interactions with databases, gateways and users. This may sound like a pure middleware play. However, today's architects are also looking for optimization of the use and control of the data, especially if more data costs more money to purchase and manage. If the goal is making the data architecture more efficient, tools exist. What may need to change, however, is the way optimal data management is envisioned.

Hugh Stewart, director, SmartStream: The individual technology components of the necessary tools and architectures to manage the growth data are available in the market, but their deployment is fractured or non-existent. The necessary technology infrastructure, software, bandwidth, data architecture designs, utility services and internal/external deployment models are all available today, but these are sometimes concealed by a whole mass of inappropriate and legacy ones. The major "tools" that are missing, which will enable the correct use of the right tools, are the "softer" elements that are associated with governance and execution of data management policy. This is sometimes caused by an absence or poor deployment of technology or poor operations processes. More often than not it is organizational inertia and anachronistic prejudice.

What needs to be done is to recognize the critical importance of data management with regard to business survival and success: trigger progressive but bold transformation initiatives which significantly reduce cost and release resource, thus freeing some of the financial savings for investment in relevant tools and services, while reallocating some of the resource savings to added value and sustaining activities such as accelerating data management governance, the timely execution of its policies, and extreme internal and external client focus.

Tony Scianna, deputy head of strategy and marketing, SunGard's Position, Risk and Operations business: No. Data has been growing exponentially for years. The difference is that now, due to what happened at Lehman Brothers and the economic crisis, regulations will require firms to access and report on more data in a consolidated, near real-time format. For example, regulators want to know what a firm's capital looks like across all legal entities, not just one. On a similar note, people are not seeing challenges around processing the data but around making it available to users.

Rather than traditional capacity tools that focus on storage platforms or networks, for example, they need technology that can aggregate data from multiple source systems from across the firm, normalize it, and consolidate it into a real-time stream that can feed any downstream system. This can help firms create a near real-time view of their firm-wide risk and counterparty exposure, as well as that of their clients. They can also use this framework to feed required data to regulators and respond to requests from compliance and senior management as they monitor risk and activities internally. This approach will also start to break down the use of silos, which prevent organizations from understanding their true exposure and risk.

Q What do you see as the top issues for data management for the rest of 2011 and 2012?

Brown: It will be knowing flow and lineage—ie macro enterprise “latency” rather than nano/micro transactional latency. Clearer definitions are needed around woolly words like “correlation” to separate identity correlation techniques from behavioral ones.

Tobin: Integrating data from different business units within our firm and being able to provide accurate consolidated management reports. For example, the same customer may do business with four different business units at Knight, but providing accurate reports across all of the business units can often be difficult because the customer may be identified differently in the various systems used in the firm.

Jacobs: I see it as latency, monitoring and the continued exponential growth in message rates, especially as options market structure continues to shift to the maker/taker market. The impact of Chicago Board Options Exchange's (CBOE's) C2 trading platform should be interesting, as that is the exchange's foray into the maker/taker model.

Brachowski: The two top issues for data management are regulation and risk. The regulatory environment in both Europe and the US is in a state of flux. Both Solvency II and the Dodd-Frank Act require an even more vigilant perspective on data management. Both of these acts are aimed at making companies assemble different types of data from across the firm into a comprehensive package to prove that they are running their business within the regulatory guidelines. In order



Robert Brachowski

Reference data product manager
Eagle Investment Systems
Tel: +1 617 219 0379
rbrachowski@eagleinvsys.com
www.eagleinvsys.com

to comply with the regulatory requirements, companies must bring together their reference data such as security master, prices etc, their portfolio-level data including trades, positions etc, and their firm-level or enterprise data such as product or sales information. Much of this data is kept in separate systems today, with differing frequencies and little to relate it directly to firm holdings.

An additional focus of the investment management industry, and thus the data management space, is risk. Companies are finding that they need to provide more data more frequently to both backward- and forward-looking risk systems. Those data requests, coupled with changes in regulation, create even more issues for companies that manage data across silos and systems. To ensure that they are meeting their stated strategies, these companies need to ensure their data infrastructure is open and flexible to handle a strategy that requires access and control of data to support the daily and monthly data requirements.

Oddiraju: We're seeing an increasing interest in “rehabbing” legacy approaches to data management for more enterprise functionality. The whole issue of enterprise data management has always been hampered by resistance to too much cost and change. Today, these issues are less meaningful, because tools for implementing an integrated data framework and enterprise data standards can be deployed without disruption to legacy systems and operations.

We're also seeing more awareness of the costs and risks associated with patchwork approaches to data management, and more sophistication about the opportunities inherent in developing enterprise data as a resource. With this approach, finding solutions to the big data issues can become more like arranging building blocks that already exist and less like re-inventing the wheel.

Stewart: The top issues are those of the main board of the firm itself; other issues must inter-relate. The key issues are:


- Cost management and reduction: bad, absent or duplicated data is a wasted cost if purchased, and toxic if part of a straight-through processing process or supply chain. Data management helps to identify, repair and sometimes convert the data to an asset that provides revenue directly or indirectly (improved service or competitiveness).
- Regulator management: to retain the license to operate is essential.

The quality, completeness and timeliness of data is at the core of an organization's relationship with regulators. Data management provides lifecycle management of this, but can also dovetail to other organizational activities to reduce the overall cost of regulation and re-leverage the processes and outputs of regulation for improved profit or use of capital.

- Client acquisition and retention: speed, comprehensiveness and quality of client-facing responses are correlated to the underlying data. STP and self-service, which is popular with clients and is cost-effective to all, relies on timely, automated processes that are very data-dependent. Speed, reliability and flexibility of handling data are client measurements that influence acquisition, retention and mandate.

- Capital and risk management: allied with cost management and regulation, the accurate calculation of capital employed/at risk and all aspects of risk management (market, credit, liquidity etc) rely completely on accurate and timely data. As they are holistic measurements, an absence of data items often invalidates measurement. There is a “Sod's Law” that the missing data is from outlying boundary events that would have a dramatic effect on a final calculation!

- Agility: an organization's need to adapt and transform in every department is vital today. Poorly managed data is always a constraint on organizational change or process improvement. One must always know your data—what it means, where it comes from, where it goes, what it does, and how it can be competitively (re)interpreted. This allows decisions on improvement, replacement or demise of function, content or process to be informed rather than hopeful.



SmartStream

Hugh Stewart

Director
SmartStream
Tel: +44 (0)207 898 0669
hugh.stewart@smartstream-stp.com
www.smartstream-stp.com

“An organization's need to adapt and transform in every department is vital today. Poorly managed data is always a constraint on organizational change or process improvement. One must always know your data—what it means, where it comes from, where it goes, what it does, and how it can be competitively (re)interpreted.” **Hugh Stewart, SmartStream**

- Core competencies: is the organization an IT or data management “shop” or consultancy? Or is it a financial services provider? What is beneficial and optimal to outsource at lower costs and better quality, and what should be managed and nurtured in-house?

Scianna: The regulatory environment is the single largest development in our industry today. As a result of the economic crisis, regulators around the world are issuing sweeping changes, not just about reporting but about how financial institutions will be expected to do business in the future. Every industry working group is responding to numerous requests for comments around the world. I have taken part in many of these initiatives and spoken on many panels and events about exactly this topic.

Collateral management, clearing of OTC trades, swap execution facilities and swap data repositories, the Office of Financial Research, new Trace requirements—the list goes on and on. Firms can't sit back and wait for the rules to be finalized; they need to prepare for these changes now. And don't forget that Dodd-Frank has over 200 rule changes, and only a handful are out for comment or are unlikely to be implemented.

Another challenge is to remove technology as a barrier to those who need to leverage data. At most organizations, data is stored in multiple systems, formats, architectures and technologies. That means that users must learn different processes and definitions. By capturing that data, aggregating it and storing it in a central location, you remove that underlying complexity and make data more accessible to anyone who requires it.

Q Are firms using a single, standardized data file framework, or is the challenge to integrate multiple frameworks into a single environment?

Brown: Neither—there are no clear strategies in this space that I have witnessed.

Tobin: I have not seen any standardization on specific report or file formats. An important part of on-boarding customers is being able to produce the reports they need. The report formats are almost always unique to the customer.

Jacobs: There are conflicting desires in the market today. On one hand a normalized format is extremely helpful for those that have used intermediaries to outsource their market data service. However, with the proliferation of unfiltered market access, many hedge funds and other market participants have become direct subscribers of market data from the exchanges.

As a result, they are used to storing data in the native format. As a service provider, the challenge is to provide a high-performance solution that can do both, and effectively reach as wide a client base as possible.

Brachowski: There is no single file framework being used for all types of data file. Firms have been able to standardize the formats of files for specific data types, such as trades or custodial data, etc. However, this format is not being used for all files. Regulatory requirements will further help solidify a standard framework used to transmit the information required by the regulations.

Some companies have implemented a standard data file framework within their organization for sending data between different systems, but they still must utilize multiple file types when sending and receiving data from other organizations.

Oddiraju: A single datafile framework is unlikely, because of the inefficiency of handling historical trending information and current relational data in the same framework.

What is possible and increasingly attractive is a fully integrated environment. The old integration paradigm, based on hand-coding that was nightmarish to maintain, once made this goal too costly. Today integration technologies not only automate coding, but offer live metadata—real-time intelligence about the data ecosystem. This enables a non-centralized, highly flexible approach to data management that is “governed” by canonical data standards and associated models, supporting anything-to-anything integration and producing ongoing new efficiencies through model-driven application development.

Stewart: Interesting words, “framework” and “environment”—they are both beautifully general and contain many sins. Quite a few firms have attempted a single data file framework, usually starting to do this in the late 1990s, particularly focused on Security Masters. Few succeeded and most failed, collapsing into multiple “silos” if they were lucky, but frequently with greater fragmentation into “potholes” with perpetual expensive maintenance. It was akin to open heart surgery, and many of the patients died on the operating table!

Therefore the reality for most companies is to integrate multiple frameworks into a single environment, which can be independent, cascade or federate, or any combination: their architecture(s) depend on the user/application requirements downstream—timeliness, accuracy, completeness, cost of data frontiers, delta change or complete etc. The world is moving to event-driven, with different businesses having different concepts of real-time, and any architectural decision has to intersect and converge on this underlying requirement.

In large, complex financial services organizations we see a pragmatic deployment with a “fabric” or “warp and weave” approach: horizontal and vertical deployments sharing and interweaving data management frameworks servicing similar needs in different departments across the organogram. This fabric approach reflects the huge efficiencies of an internal utility model, the virtues of which can be extended towards an industry utility view of wider virtual fabric.

Scianna: Each asset class has highly imbedded solutions with nuanced processing requirements. These solutions, whether they have been bought off-the-shelf or built in-house, are a barrier to consolidating on a single data framework.

The idea of standardizing on a single data framework also introduces numerous questions around feasibility. Which single data framework is equipped to operationally support the process requirements of equities, fixed income, derivatives and FX products across the globe?

Instead, firms can retain their existing, individual systems while enjoying the benefits of a single framework. By pushing that data into a normalized and standardized environment, they can feed a single stream of data into any downstream system. Bringing disparate data into a single holistic environment helps firms handle regulatory and market changes.

Q Should these issues be addressed from the individual business lines or centrally from within the enterprise?

Brown: It is an enterprise issue. But most enterprise central functions are either oblivious to the problem or lack sufficient domain expertise to create a useful solution to the line of businesses, and end up being seen as wasteful, obstructive bureaucrats.

Tobin: In general, all of our transactional data is pulled into a single data warehouse. This gives us the ability to produce consolidated reports across the firm. We do have some data warehouses that are maintained by specific business units, often containing market data and transactional data that is analyzed to improve the performance on algorithms and smart routers.

Jacobs: These issues should absolutely be addressed centrally. While there may be an intentional decision to silo certain market data solutions, the ramifications of this choice should be shared. Again, the challenge of raw, non-normalized data versus a normalized format has enterprise-wide ramifications, especially as market data is integrated into complex-event processing (CEP) systems, algorithmic trading solutions, for back-testing and simulation, for regression testing, and, of course, for trading. Different needs may require different formats, and a technology choice in one area could have substantial downstream ramifications.



John Jacobs
Lime Brokerage

Brachowski: Data management solutions should be centrally located within the enterprise. However, they must be built with input from all the business lines in the organization. While this allows companies to gain the traditional benefits of centralization, such as improved consistency, reduced redundancy, and so on, relying on the business lines to drive the data management solution will help ensure that the business process and the data management solution stay aligned.



Volanté 
The Messaging
Company

Vijay Oddiraju

CEO
Volanté Technologies Inc
Tel: +1 212 905 6231
vijay@volantetech.com
www.volantetech.com

Oddiraju: The reality is that budgets are rarely dedicated to strategic enterprise initiatives. Budgets are more likely to be dedicated to a tactical, localized business purpose. Unless there is a clear case for payback, support is unlikely to emerge at the senior level.

Right now, that understanding has been rare, at least in terms of enterprise integration strategy, because it is a relatively new concept. Fortunately, it is a technological approach that quickly delivers ROI on the project level. And except for some IT departments or consulting firms, where handcoding may be regarded as a kind of employment insurance, the technology doesn't trigger political issues. But even if that were not so, the intense pressures on the data architecture from regulatory and business demands are real. As a result, the strategic implications of integration choices—especially in the impact on performance, efficiency and cost of ownership—are becoming clearer.

Stewart: The answer is both, or why have a “joined-up” corporation? Also, why have a choice of two: there is a third option too of an external outsource or utility contributor. The triangular model of governance, ownership and execution determines how the question is answered. Governance and policies should be central. Ownership of data is distributed throughout the organization and thus is business line driven. So far it is easy to be deterministic and Stalinist. Execution of the policies of governance and their deployment, however, is subtler and the binary approach inefficient. As mentioned earlier, an interwoven-fabric approach has to be adopted where different content, processes and functions can be managed by a mixture of individual lines who might also act on behalf of other business lines and also might cascade their results and outcomes to the center and

central departments who in turn service business lines. The fabric is determined by an “efficient frontier” of cost, pragmatism and business/data management drivers as mentioned earlier—like the current arguments for the bidirectional generation of electricity, which can come from both a central utility and also from homes and businesses, from the excess of their internally-generated capacity. Thus a balanced commercial and operational ecosystem is created.

The third option is to extend the fabric with the dimension of an external outsource or utility contributor partner who further improves the “efficient frontier” of cost, quality, reliability, technology and timeliness.

Scianna: It really depends upon how a business is structured and the functions that need to consume the data. There is no one-size-fits-all solution or approach that will work for every organization. However, working in isolation within individual business lines can prevent senior management from fully understanding enterprise risk and exposure. A centralized approach is most suitable for helping firms prepare for regulatory change, understand their activity, and support business growth.

Q What role has unstructured data taken on for most firms and how does its management differ from structured data?

Brown: Few firms are addressing the problem strategically, despite significant improvements in tooling to analyze, index/catalog and manage flow/delivery/integration/reuse. Generally the approach is reactive such as to regulation/audit/storage/retention costs rather than proactive use and leverage.

Tobin: Almost all of the data we store and manage is in a structured format, with the exception of news feeds which are stored and indexed for retrieval by users, traders and customers for reading.

Jacobs: Unstructured market data is still in the early stages of deployment. Just this week [in mid-May], Derwent Capital Markets, a London-based fund, will be monitoring Twitter as a market-sentiment guide to aid in its trading decisions. Although the opportunity appears to be large, the benefit has yet to be explicitly and quantifiably defined. Therefore, there does not appear to yet be a consistent set of

“

“Data management solutions should be centrally located within the enterprise. However, they must be built with influence from all the business lines in the organization.” **Robert Brachowski, Eagle Investment Systems**

“

“Among our customer base, we see the desire to greatly improve internal efficiencies in handling data in and out of Excel, MS Word, and ODF formats for reporting purposes, generating statements, and for the automation of message data.” Vijay Oddiraju, Volante

packaged solutions, or a standardized approach on how to manage and implement. This area is still in its infancy, despite the fact that a 2008 report from industry consultancy Aite Group forecast \$141 million in spending on unstructured data in 2011.

Brachowski: The proliferation of data means that more and more data is being arranged in an unstructured format. Unstructured data is taking a larger role at financial firms than ever before. Companies are starting to understand that structured data is the essence of their business, while unstructured data influences their business. The biggest concern when dealing with unstructured data is learning how to properly identify it and then link it to the structured data that the firm already manages. In addition to being able to relate and store unstructured data, the most important aspect of unstructured data is the ability to search it and allow it to be presentable when required. Unstructured data cannot be managed in the same way as structured data. Most companies are using tags or metadata elements to organize and relate it back to their fixed structures.

Oddiraju: Among our customer base, we see the desire to greatly improve internal efficiencies in handling data in and out of Microsoft Excel, Word and ODF (Open Document Format) formats for reporting purposes, generating statements, and for the automation of message data. Some of these techniques for tagging and extracting key data are now being applied to automatically generate payment messages and bank-to-bank financial exchanges, as well as for customer or counterparty statements of accounts or assets.

Stewart: Unstructured data has existed forever for organizations and there have been many revolutions in its technology and management, such as printing, keyword search, Google, pattern-matching logic etc. We are going through another revolution as the weaknesses of individual human and social intuitions, as well as abilities to assess risk naturally, are being revealed as being misleading and downright stupid. Our cave-person brains, bodies and derived natural, social behavior are seen to be inappropriate to the new complex, quantum, information-intense world. Now there is a close correlation between how an organization benchmarks with their peers in the way that unstructured data is corralled, made manageable, exploited, and reliable insights thus gained that can be acted on.

The letters “BI” now are appended to many people who are also working in the conventional reference and market data spaces as their remit and responsibilities expand to accommodate the business imperatives.

Unstructured data, at first sight, should be categorized and managed separately from market and reference data. But just as market and reference data have merged, overlapped and “bled” into each other, so too have structured and unstructured data. Nowadays the disciplines inter-mingle. Fuzzy logic pattern matching helps when it comes to cleaning up entity databases and searching for missing information; or analyzing tick databases to identify trading strategies or—with adaptive neural networks—derive new strategies. It can also help with interpretative screen-scraping for data capture and acquisition in order to replace manual keying in.

Even good old reference data management can use the unstructured data from the exception-reported events from the cleaning of data vendor files to predict anomalies and distortions of sources as we do at DClear which increases our data quality metrics. Even good old reference data management can use the unstructured data from the exception reported events from cleaning data vendor files to predict anomalies and distortions as we do at DClear.

So unstructured data should also fall into the triangle of data management roles and responsibilities—governance, ownership and execution; avoiding the silo inefficiencies but adapting to a fabric architecture; allowing the dedicated technologies, processes and methodologies of reference, market and unstructured data to bleed and merge into each other. Finally, this radical widening of the footprint of the responsibilities of data management departments also strengthens the need to identify internal and external utility deployment roles, as there are no alternatives if one is to keep up covering the geometrically expanding ground.



SUNGARD®

Tony Scianna

Deputy head of strategy and marketing
SunGard's Position, Risk and Operations
business
Tel: +1 646 445 1073
anthony.scianna@sungard.com
www.sungard.com/streamforposttrade

Scianna: In spite of efforts to modernize data sets and standardize on industry protocols, at most firms a great deal of data still remains outside the logical framework. Many organizations have not yet embraced the unstructured data revolution because the technology is still maturing. There are many aspects of the technology that are very appealing, but it is still unproven. ■



Are you confident about your **DATA MANAGEMENT STRATEGY?**

Global regulatory changes are creating a series of business and technology challenges. The need to consolidate your data across the enterprise into a single view has never been more important than it is today. Does your data management strategy reflect the post-crisis changes and the expected regulatory requirements? Is it flexible enough to remain effective even as regulations change?

SunGard can help.

Learn how **StreamCore®** can provide a holistic view of your organization and your clients for effective and flexible risk and exposure management.

www.sungard.com/streamforposttrade/waters

waters

