

ConteX MCS 4.0

Design criteria and product enhancements

– a Gissing white paper

Table of Contents

| | |
|----------------------------------|---|
| Introduction | 3 |
| Changing Contributions Landscape | 3 |
| Multiple Platform Support | 3 |
| Functional Requirements | 3 |
| Design Criteria | 4 |
| Improved Performance | 4 |
| Performance Results | 5 |
| New Functionality | 5 |
| Improved Administration | 6 |
| Future Enhancements | 6 |
| Conclusion | 6 |

Introduction

ConteX MCS R3 has become, since its release in 2002, the de-facto standard for multi-vendor contribution systems in investment banks and major inter-dealer brokers worldwide.

At the time of writing, 10 of the world's top 15 banks use ConteX MCS to manage some or all of their multi-vendor market data contributions. Indeed we have evidence that ConteX MCS is being used as a performance benchmark by which developers of similar systems measure themselves.

Changing Contributions Landscape

However the world of contributions has changed and evolved since the release of ConteX MCS R3 in a number of ways:

- Increased number of instruments
- Increased rate of price changes
- Extension of contribution to additional instrument classes
- Higher communications data rates
- Growth in algorithmic trading

When we reviewed the impact of all of these changes taken together, the design team at Gissing came to the conclusion that a radical redesign of the architecture of ConteX MCS was necessary. We felt that we could not only keep up with the massive increases in data volumes, but also start to anticipate future requirements.

There have been some developments in both operating systems and hardware which help with this:

- Significantly higher processor speeds
- Multi-core CPUs

Indeed at the time of writing (May 2009) systems with dual quad-core CPUs – effectively 8-processor Intel® systems – are available for around £3,000 (\$6,000). This indicates that the price performance of available systems is improving massively. It also means that we at Gissing can continue with our strategy of developing for the Intel/Microsoft® Windows® platform without compromising system performance.

Multiple Platform Support

A handful of clients and prospects have mentioned that they would like to see ConteX MCS available on one flavour of UNIX® or another. However most of them have licensed ConteX MCS on Windows anyway.

The few requests we have received for a UNIX-like operating system version of ConteX MCS concentrate either on Linux® (Red Hat® and SUSE™ are the main distributions in commercial use) or Solaris for Intel or Sparc. It is interesting to note that there are many fewer flavours of UNIX in widespread use today than there were even five years ago. We are keeping this requirement under review, but there are currently no plans for any UNIX version of ConteX MCS.

Functional Requirements

We began the design of the next version of ConteX, to be numbered 4.0, during 2006, and had solicited many views from existing users and clients regarding desirable features and enhancements. These primarily included:

- Reduced latency particularly to improve suitability for use with algorithmic and programmatic trading applications
- Improved performance to manage larger numbers of contributed instruments (one European client expected to increase the number of equity derivatives having prices contributed from 10,000 to 60,000 instruments by the end of 2007).
- Improved performance to manage very large increases in volumes of contributed data (the same client expects their current rate of contributed prices for equity derivatives to go from 10,000 prices per second to 50,000 prices per second in the same time).
- Easier, and remote, administration of the system
- Better permissioning and partitioning of single ConteX systems
- Dynamic reconfiguration after instrument changes and additions

It has also become clear to us that along with the massive increase in contributions volume, there are some new requirements for a firm's contribution system to be able to manipulate data while "in flight". This certainly doesn't mean building an entire pricing engine or pricing application into the contribution system, but it does entail providing capabilities to perform relatively simple arithmetic and string manipulation on contributed data in order to derive new data items for onward publication.

A further requirement has been identified, primarily as a result of the provisions of the European "Markets in Financial Instruments Directive" (MiFID). Gissing has been focussed primarily on the transparency provisions of MiFID, which came into force across Europe on November 1st 2007. These require some organisations (Systematic Internalisers –SIs) to publish

pre-trade quotes in real time to a public venue (there are some circumstances in which this is not required, but it's beyond the scope of this white paper to cover all the details). The provisions also require all trading organisations to ensure that the results of the trade are published in near real-time – and in any event within three minutes of the trade.

Conceptually there is nothing complex about using an existing contributions system to meet these requirements, provided the source data is available within the appropriate time limits. However there is a further requirement within MiFID that data contributed to a suitable venue to meet these requirements is validated outside the trading system on which the trade was made, prior to being made public by the venue.

Some venues (Project BOAT for example) provide validation as part of their contribution protocol, however many do not, so we decided to introduce a new validation capability within Contex MCS to enable contributors to validate contributed data, both for MiFID and for regular publication and trading applications, prior to its onward transmission.

Design Criteria

There are many definitions of latency, and as many ways of measuring performance of a system. However we decided that we needed to address every aspect of the path of data through Contex MCS, including:

- start up and subscription time,
- the time data spends within an input handler,
- the time it takes for data to be routed from input handler to output handler
- the time data spends within output handler(s),
- the efficient contribution of data to a destination, including conflation and load balancing between multiple connections, and
- the number of processor cycles required to manage, route, manipulate and contribute the data.

Addressing all of these issues necessitated a radical redesign of the entire Contex architecture and particularly the inter process communication (IPC) between components of Contex.

In order to do this, we decided that there were some design criteria that should be adopted:

- To avoid, wherever possible, allocation of system memory in real-time. Allocation of memory is highly detrimental to performance; it is expensive in processor cycles and can cause contention between threads so compromising parallel processing on multi-core and multi-processor systems

- To totally separate data from status and control information – so that system status is always available and the system is available to, and responsive to, the administrator at all times, irrespective of the data load at the time
- To avoid Windows message queues and any other technique that has any risk of overflowing in a high-load situation. Enabling Contex to always handle the message load will improve both stability and overall system performance
- To rationalise the queuing of messages in the Kernel. In the existing version of Contex, messages passed through several Kernel-managed queues, including a queue for each thread of communications. In the new architecture this mechanism has been substantially simplified. This, combined with intelligent queue management, significantly reduces both processing requirements and any risk of queuing delays.

Improved Performance

When starting a Contex system, the system needs to subscribe to all the necessary data. Historically Contex has sent a separate message for each data item that it wants to subscribe to. In systems with large numbers of instruments this could entail many thousands of messages and take some time. Contex MCS 4.0 has been designed with an "Active Configuration" which introduces the concept of Auto-Subscription. This means that on start up a single message is generated which implicitly opens all the necessary message streams. This also removes the necessity to have subscription logic in the Contex Kernel, so improving its performance. Combined with the separate streams for status and control information the system administrator will always be aware of the status of Contex as it starts up.

Active Configuration also means that Contex is able to pre-calculate much of the routing information for data before it arrives. This greatly streamlines the processing and routing of data – it avoids having to subscribe to record level data that's not actually required, so this releases processor cycles that would previously have been used to process and queue this data before discarding it as unnecessary.

To further improve message routing and processing efficiency, Contex MCS 4.0 introduces the concept of "Tear-Off Messaging". In previous versions of Contex MCS, the Kernel had to decode each message in order to be able to divide it into its constituent parts and route the appropriate fields to the right places. This consumes a great deal of machine resource and time to decode and interpret large parts of a message that actually aren't needed anywhere else.

Introduced in Contex MCS 4.0, Tear-off Messaging is a technique where the Kernel can identify the useful parts of a message, and route them to the appropriate destinations, without having to decode the message itself or understand its content. Only the components of the message that are needed by each output handler are routed to it.

This leads to substantial performance improvements because the Kernel doesn't need to process any of the content of the message. The handlers are also more efficient because they do not have to receive and process parts of a message they don't need, reducing the internal bandwidth required. Furthermore, new messages to be used by internal processes can easily be added to the system without having to modify the Kernel, because it doesn't need to be able to decode and process the message content.

Performance Results

Clearly each configuration of Contex MCS is different. Performance of a Contex MCS system will vary depending on a vast number of parameters including the configuration of the server on which it's running, the precise configuration of input and output handlers, the actual reformatting of data being undertaken within Contex, the system loading, the number of instruments and updates and the precise format of the messages.

However in order to provide some indication of the improvement of performance, we have undertaken some side-by-side comparisons of Contex MCS 4.0 with its immediate predecessor.

In these tests we have been able to demonstrate that end-to-end latency for a single message has been reduced by approximately an order of magnitude, to a few hundred microseconds or less.

We have also been able to show that where previously a fully active Input handler might have taken up to 25-30% of CPU cycles at peak, in Contex 4.0 the functionally identical input handler processing the same data now peaks at around 6% of CPU. This clearly means that more processes can be run on the same hardware, and therefore a greater number of contributions can be handled simultaneously.

New Functionality

With Contex MCS 4.0, Gissing has introduced its own expression language, Gissing eXpresso™. This is similar in many ways to the formulae used in Microsoft Excel® so will be immediately familiar to many users. The language has over 40 functions, including both

mathematical and string manipulation functions. It also has simple conditional logic (if, then else) capability.

To ensure maximum performance Gissing eXpresso functions are compiled at entry time, so the functions execute at machine speed – rather than as an interpreted language compiled at run-time.

This language is used in two new, optional, components of Contex MCS 4.0.

The first of these is the Derived Data module. Although Contex MCS R3 provided limited data manipulation functions such as the ability to scale numbers, Gissing eXpresso provides a much richer set of functionality which enables derivation of data while the message is being processed. Simple examples include being able to derive the instrument name for each destination from the input record, or being able to derive a "mid" or "spread" from the "bid" and "ask" prices in the input record.

It should always be borne in mind, however, that the more complex the data derivation, the more processor cycles it will consume and this means a penalty in increased latency. We recommend that clients experiment with different formulae and expressions while monitoring system performance, throughput and latency to achieve their preferred combination of high-performance low-latency contributions and the necessary data derivation they need.

The second new component of Contex MCS 4.0 is the Data Validation module. This uses Gissing eXpresso to establish one or more validation criteria for input data. These criteria can now include historic data – Contex MCS 4.0 is now capable of storing an administrator-defined number of previous values of a field in a record. These can be used to calculate moving averages, or high and low values for a price, for instance. Then if the newly contributed value falls outside the defined range it can be rejected.

The behaviour on rejection can also be determined by the administrator. It can cause an error but still be contributed, or it can clear the record to avoid either incorrect or stale data being displayed, and optionally an error can be returned to the submitting process.

Improved Administration

Many of the comments and suggestions from existing users of Contex MCS involved the management and administration of the Contex system. In Contex MCS 4.0 we have substantially improved this area of the product and have included:

- Remote administration from any Windows PC on the same network segment via a simple TCP/IP connection.
- Improved alerting
- Dynamic application of configuration changes. This means that new instruments, data derivations and validation can be added and applied immediately – without having to restart (or “bounce”) the server.
- Scheduled application of configuration changes. As above, but the changes can be scheduled for implementation at some later time.
- **Statistics and Diagnostics:** The improved user interface in ConteX MCS 4.0 allows the collection and presentation of a whole range of statistics not previously available. This includes being able to view the individual latency for a particular record as it passes through the system – both the overall “end-to-end” latency and the time it spent in each component of the system. Average latency is also available in the same way. This enables much easier tracing and diagnostics.
- **Status Instrumentation:** In ConteX MCS R3 status information was published in pages intended to be human-readable. ConteX MCS 4.0 replaces this mechanism with a data feed. Not only can this be directed into a third-party status tool such as IBM® Tivoli® or a Simple Network Management Protocol (SNMP) tool, but it also substantially reduces both the CPU requirement and the verbosity of the information being published. Furthermore much richer and more descriptive data can be provided, including more useful attributes, such as classifying errors as “serious” or “recoverable”.

ConteX system is being used to manage contributions from multiple desks or departments – each desk can be given the rights to manage their own contributions while ensuring that there’s no risk that they inadvertently change something to do with another desk’s activities.

Conclusion

ConteX MCS 4.0 is undoubtedly the highest performance, lowest latency version of ConteX MCS yet developed. It has been specifically designed to meet the requirements of high volume contributions in today’s dynamic trading environment, while it’s also easier to use, manage and configure than before.

It also adds a range of new functionality for manipulating, deriving and validating data in real-time, and provides a platform for some exciting new features and functionality for the future.

ConteX 4.0 raises the bar for multi-user contribution systems.

Founded as Gissing Consulting in 1991, Gissing is now the leading supplier of multi-vendor contribution systems to financial institutions worldwide. We have more than 30 major investment banks and brokers as clients and offices in London, New York and Hong Kong. Gissing was acquired by Reuters in January 2008 and now operates as a division of Thomson Reuters.

Future Enhancements

The new internal architecture of ConteX MCS 4.0 has also been designed to enable a number of future enhancements including:

- Further improving user interfaces and status information, including the addition of data visualisation tools.
- Greater scalability, including possible support for ASP-like hosted versions of ConteX MCS.
- Extension of derived data to include parameters or data from other input records or third-party sources such as reference databases.
- User configuration of Kernel threading strategies to enable tuning to different hardware configurations.
- **Partitioning and improved permissioning:** ConteX MCS will allow the establishment of different levels of user permission – users can be allowed to view but not to change anything, and can also be restricted to one area of ConteX, such as a particular set of instruments. This will be particularly useful where a