



The CommSecure Markets Information System

*Delivering online transactions with
new levels of performance, reliability
and cost-effectiveness*

**A CommSecure white paper
March 2002**

This White Paper is the property of CommSecure Limited (the "Company") and the information contained herein is protected by copyright. No part of this white paper may be reproduced in any form including photocopying, electronically or translation into a foreign language without the express written permission of the Company. Such unauthorised usage is a violation of Australian and International copyright law.

The Company considers the information contained in this White Paper to be reliable but it is not warranted by the Company. The document may also contain opinions of the Company, which are subject to change from time to time.



Solving Online Challenges

Businesses face significant challenges in electronic commerce as they seek to increase the volume, speed and complexity of online transactions. They must ensure customers enjoy a high quality of service without compromising overall corporate objectives – particularly revenue growth and cost reduction.

This is especially the case in the delivery of web-based financial services, where key offerings such as share trading, online banking, news content and market updates must be provided quickly, reliably and securely to a wide variety of customers.

As customer demands increase, businesses are coming under greater strain. There are three principal sources of pressure:

- **The cost of supporting millions of transactions.** These costs continue to rise as businesses make regular changes to old technology platforms, often extending systems beyond their original scope.
- **The challenge of generating new revenue.** Many businesses are unable to introduce new services on existing systems, restricting their ability to offer new products and increase revenue.
- **The risk of a system outage.** This is a constant danger in online finance, raising the prospect of technical difficulties that undermine customer satisfaction, harm ongoing customer loyalty and lead to negative media coverage.

The explosive growth of the Internet in recent years has contributed to today's challenges. The current generation of transaction processing technology was designed for an earlier era and does not position companies to take advantage of future opportunities. Unfortunately, business goals are being limited by technical capabilities.

It is now clear that Internet platforms intended for the first phase of online services are not appropriate for a new generation of more complex services with more intensive technology requirements.

Businesses have the potential to deploy a new generation of technology to address these issues. Because it is based on a new design, the CommSecure Markets Information System addresses today's key challenges of cost reduction, revenue growth and system reliability. The system has been commercially proven in financial organisations with highly intensive transaction loads.

Technology barriers

As business leaders seek to achieve their most important corporate objectives, they face several technical barriers that tend to stymie their efforts. These technical issues currently limit the potential to improve the capacity and efficiency of online financial applications.

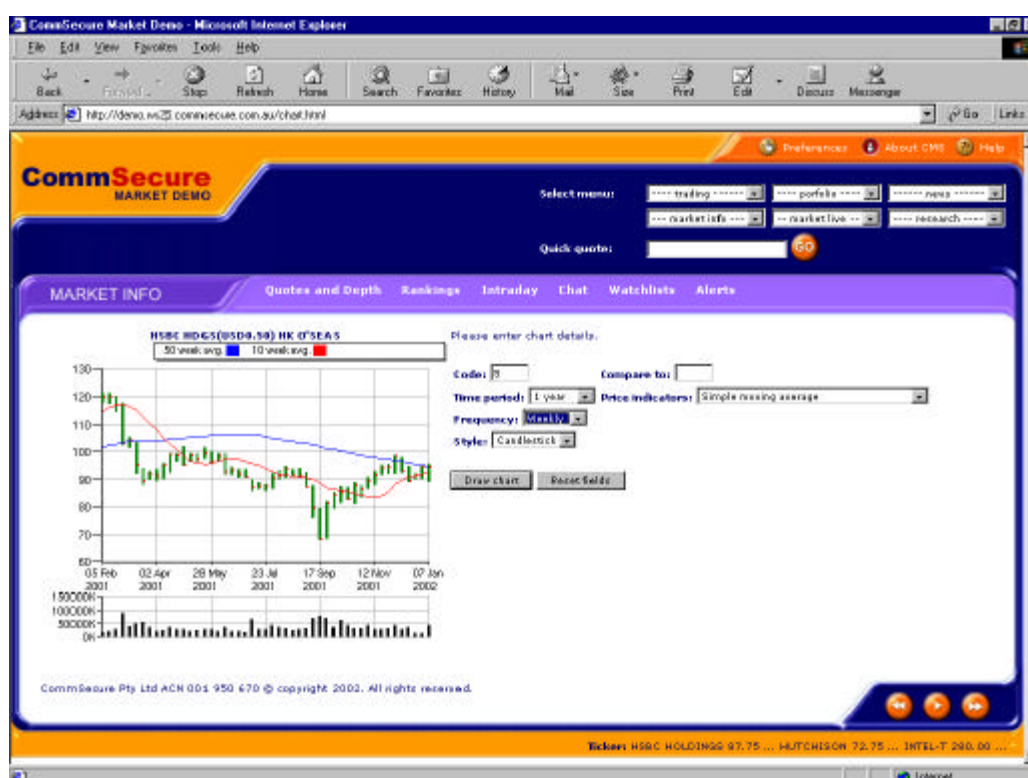
Organisations are required to provide large amounts of disparate data from various sources in order to satisfy customer demands. This data is presented using multiple channels including interactive voice recognition (IVR), Internet, wireless, and custom applications. These data formats cover a broad spectrum of information ranging from local content – such as information from a local stock exchange and news services – to global content – such as major indices and world news.



In the majority of cases, the most critical demand is for local content. This makes it vital to deliver local information in the most efficient manner. The challenge is to extract the data directly from the originating source and provide a mechanism to deliver the data as efficiently as possible.

Although customers also seek global content, these demands are generally less urgent. Understanding the demand mix between critical and non-critical content is an important factor in designing a system that ensures customer satisfaction.

Adding to these issues is the extraordinary fluctuation in transaction loads. The financial markets information system must meet the demands of a rapidly shifting financial marketplace. On volatile days, usage during opening and closing trading periods can see the load on the system increase to a level many times higher than normal usage. The ability to handle a rapidly escalating load on the system is also a critical issue.



If properly designed, an efficient financial markets information system will lead to an increase in customer numbers at the same time that individual usage grows. This places even greater pressure on the system to sustain robust performance.

The ultimate challenge is not only to implement a system that performs efficiently, but to have incorporate the ability to inexpensively scale the system rapidly upwards in order to achieve industry-leading performance as customer demands grow.

A new platform

The CommSecure Markets Information System is a financial markets technology platform that provides the essential infrastructure for organisations to deliver real-time financial markets information to multiple presentation channels.

This new platform solves the logistical problems inherent in the mass delivery of real-time data. In doing so, it removes the need for customer channel software to deal with logistical issues.



The CommSecure Markets Information System delivers improved throughput to existing customer channel applications with the ability to scale capacity upwards to keep abreast of and exceed the requirements of a rapidly growing customer base.

The CommSecure Markets Information System is a high performance system that provides online real-time equities trading information. Using a single interface to this system, businesses can provide their financial customers with immediate access to content including:

- live stock quotes
- delayed stock quotes
- historical, intra-day and analytical charts
- news
- market depth (bid/ask queues)
- world indices
- research
- futures and options

The system can connect to multiple stock exchanges in a variety of languages and currency representations. It can also be customised to connect to any data sources required.

Because the CommSecure Markets Information System is designed as a middleware technology platform, it enables the smooth flow of data from any other applications and financial systems. It facilitates the transmission of data from data suppliers (such as stock exchanges and data feeds), processing the data into content and then delivering the content to a presentation channels. These channels include, but are not restricted to, formats such as websites, IVR systems, the Wireless Application Protocol (WAP) and other wireless applications.

Key benefits

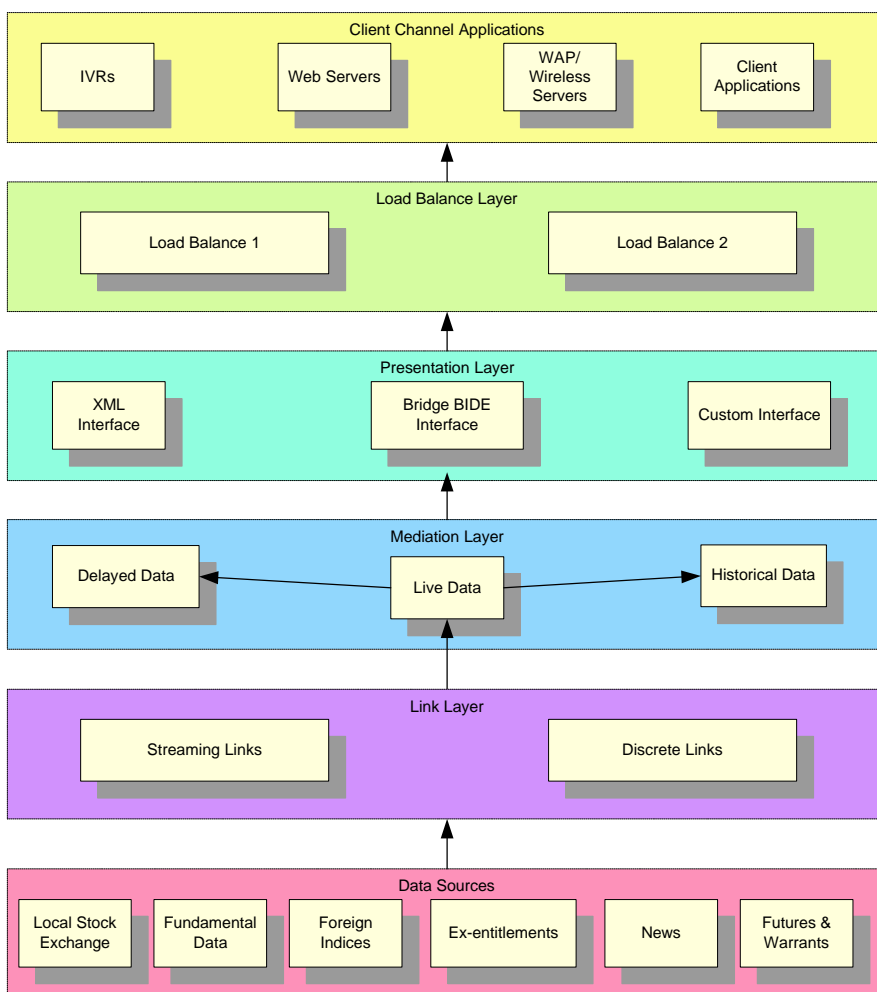
By overcoming significant technical hurdles, the CommSecure Markets Information System brings major benefits to financial organisations:

- **Reliability.** The system achieves high availability using a fault-tolerant design with very few single points of failure. Using duplicate server sites, the system has automatic fail-over capabilities to maintain operation during outages.
- **Scalability.** The system can be expanded in a straightforward and cost-effective way to meet demand during peak usage. The system is monitored regularly and upgraded if necessary to ensure rapid response time and customer satisfaction.
- **Speed.** The system is built to achieve high transaction speeds so that it can maintain performance levels during peak trading periods on volatile days, without degradation in performance.
- **Open standards.** The interface to the system is designed using open standards that ensure easy access, principally using the extensible markup language (XML) or the BIDE interface.
- **Flexibility.** The system's presentation layer uses customisable parameters that enable businesses to determine the visual aspect of their content. The separation of dynamic content from the visual presentation means that corporate branding can be maintained for each client.



System Overview

The CommSecure Markets Information System is a distributed system consisting of layers of communicating hosts. Each layer consists of multiple hosts, and hosts can be added to each layer to increase capacity. Hosts in each layer request data from hosts in the lower layer using TCP/IP or via direct connections to databases. Streaming data sources feed the system using TCP/IP and a variety of data source specific protocols, and discrete data is received using HTTP, FTP, SSH mechanisms or whatever data sources require. Client applications use TCP/IP to request services from the system and additional custom interfaces are made as required. Multiple client protocols are supported.



The CommSecure Markets Information System supplies data to multiple applications or instances of customer channel applications (top of diagram) with data from multiple data providers (bottom of diagram). The design of the system harnesses the power of distributed processing to spread processing load, remove bottlenecks, and transparently handle component failure



The CommSecure Markets Information System consists of four layers:

The Link Layer components provide data to Mediation Layer components. There is at least one Link Layer component for each data source. Streaming Link Layer components present a CORBA interface to clients. Discrete Link Layer components obtain files which are copied by Mediation Layer components using SSH.

Streaming Link Layer components manage dual redundant links. Each streaming component manages a single link. Only one link is active at a time. The inactive streaming component obtains a data stream from the active streaming component in order to supply data to its Mediation Layer clients. Up to ten streaming components can use another streaming component as a data source. A component that uses another component as a data source can itself be used as a data source in order to scale capacity upwards.

The Mediation Layer is replicated on a number of hosts to achieve the required performance capacity. There is at least one set of Mediation Layer components shown above for each streaming Link Layer component. A streaming link layer component can support up to four sets of Mediation Layer components. Live Mediation Layer components poll Link Layer components and write to the live database. A snapshot of the live database is periodically sent to the delayed database by the Delayed Mediation component and the historical database is updated after each trading day by the Historical Mediation component. The admin database contains user and password information and is updated by HTTPS.

Mediation Layer components randomly choose a route to a Link Layer component and make a connection to it. If there are two possible components that support the desired service and two LANs there are four possible paths from which to choose. Requests are sent via connection until an error occurs. If an error occurs a new random connection is made.

The Presentation Layer is replicated on a number of hosts to achieve performance capacity. That is, there are many Listeners. Listeners listen for requests arriving from either LAN, accept the incoming request, determine the interface the request corresponds to, and call the appropriate interface object to process the request. Listener components contain a pool of database connections. The pool is initialised by connecting to each Mediation Layer database via each LAN. A random database connection is used to service the potential request from the pool. If the connection fails it is removed from the pool. When the pool is empty a new pool is generated.

The Load Balance Layer - Client Channel Applications send requests to an IP address that is shared between two load balance machines. Only one load balancer receives requests at any given time. The inactive load balancer pings the incumbent in order to determine its status. If the incumbent fails the inactive load balancer will claim the IP address. If both claim the IP address only one will succeed and the other will become inactive.

The incumbent load balancer receives requests from clients and sends them on to a Presentation Layer machine network interface randomly chosen from a list of target Presentation Layer machine IP addresses (there are two per machine). Load balancer machines poll Presentation Layer machines on a continuing basis using a ping request. If Presentation Layer machine network interface fails it is removed from the list of target IP addresses.



The inherent design of the CommSecure Markets Information System solves capacity problems in two ways. First, data path bottlenecks are replaced with a high capacity service. As usage increases, the system is upgraded to increase the capacity of the data paths used by customer channel applications.

Second, CommSecure Markets Information System facilitates removal of bottlenecks in customer channel applications. The system is interposed between customer channel applications and the resources they share. By removing direct access to shared resources, the CommSecure Markets Information System allows existing applications to be replicated and run concurrently to increase capacity to users. Removing access to shared resources has an additional advantage for customer channel applications – they become simpler and easier to maintain.



System Architecture

The CommSecure Markets Information System architecture uses a multi-tiered N-redundant client/server architecture on a dual redundant network to achieve its scaling objectives. Tiers correspond to the layers discussed already. Tiers contain multiple hosts in order to provide computational capacity and redundancy. A 'client/N-server' mechanism coordinates between components within in adjoining tiers to harness the computational capacity.

Component Replication

Components in the Load Balance Layer are replicated for redundancy. There are two load balance hosts. The incumbent load balancer accepts incoming connections and forwards them to Presentation Layer hosts. The aspirant load balancer is a backup that takes over if the incumbent load balancer fails.

Components in the Presentation Layer are replicated across multiple redundant hosts in order to supply increased computational capacity. Presentation Layer components make database connections to hosts in the Mediation Layer. The database connections have significant latency so multiple Presentation Layer components run on each Presentation Layer host to more efficiently utilise CPU capacity.

Components in the Mediation Layer are also replicated across multiple redundant hosts in order to supply increased computational capacity. Mediation layer hosts have two hard disks to reduce disk head movement latency during database access.

Link Layer components are replicated in order to provide redundancy for control of streaming data sources. Streaming data sources such as a direct stock exchange link tend to be point-to-point connections managed by a Link Layer host. If the host fails, a backup host recovers and restores and takes over processing. Other streaming data sources are broadcast servers. Link Layer components that are clients of these data sources are distributed among Mediation Layer hosts to reduce the length of the data paths.

Client/N-Server Mechanism

Component relationships between the Load Balance, Presentation, Mediation and Link layers have strictly tiered client server relationships. Components in the Load Balance Layer only invoke components in the Presentation Layer. Components in the Presentation Layer only invoke components in the Link Layer and so on.

Each component is a client that can access N servers. A Load Balance component can connect requesters to one of a number of Presentation Layer Listeners. A Presentation Layer component can access databases from any Mediation Layer host. A Mediation Layer can access one of a number of Link Layer servers.

A client component connects to a server by randomly choosing between the data paths available. The available data paths consist of the routes to each of the hosts in the adjoining lower tier that can be reached by each of the LANs. Connection failures result in the removal of data paths from the list of available data paths. This mechanism is used by components in the Load Balance, Presentation and Mediation Layers to access components in the next layer down.

Dual Redundant Networking

Communication between components on different hosts is via dual redundant LANs. Hosts have two Ethernet cards.



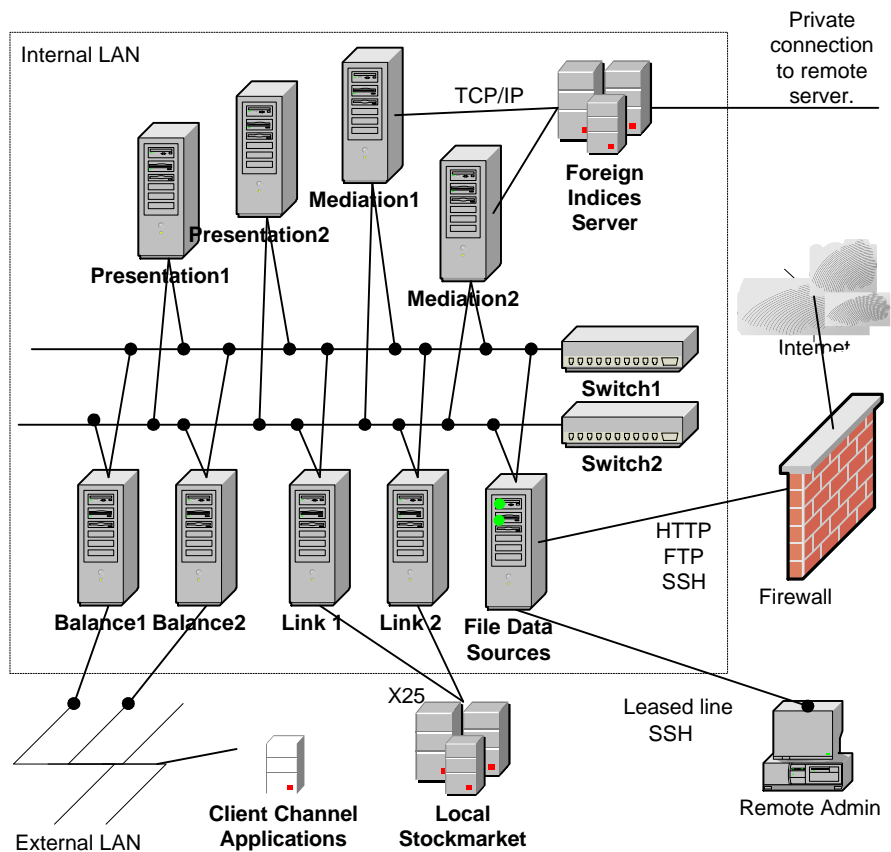
The backbone of the architecture is a dual redundant 100Mb switched LAN. Hosts are connected to both LANs.

The Load Balance Layer always consists of two hosts. Hosts connect to multiple client LANs. Load Balance Hosts have four Ethernet cards.

Configuration Example

The number of machines for a particular installation may vary and additional machines may be required, depending on the desired throughput and response times. The diagram on the next page shows the implementation installed in HSBC Hong Kong. This configuration used nine 800MHz dual Pentium III machines. The HSBC installation uses a dual redundant system to achieve a response time of 12ms with a throughput of 2.5M requests per hour and a peak of 320 requests per second in production. Before going into production the system was benchmarked and was shown to sustain 500 requests per second with an average response time of 25ms.

The pre-implementation benchmarking also addressed the scaling of the system. It was shown that by using 20 presentation machines, 20 mediation machines and 20 link machines, the CommSecure Markets Information System could sustain this response time at 20M requests per hour.



In the dual-redundant configuration there are two presentation machines, two mediation machines and two link machines. These are 800MHz dual Pentium III machines. Mediation machines have 2GB of RAM, and link and presentation machines have 1GB of RAM. There are always two load balance machines.



Mediation1 and Mediation2 each contain a full set of Mediation Layer components. Link1 and Link2 each contain a single streaming link component. The File Data Sources machine houses the discrete link components. Presentation1 and Presentation2 each contain 40 Listener components.

The streaming foreign indices are provided by a third party from a remote site. The streaming foreign indices link component is provided by the third party and resides on the Foreign Indices Server. The Foreign Indices Server has a private connection to the remote site.

The File Data Sources machine has a connection to the Internet to obtain file data. No data is served to clients through this connection. The Link Layer components on this machine use whatever security protocols are required by the servers from which data is procured, as long as this cannot compromise CommSecure Markets Information System – hence the use of insecure protocols such as HTTP and FTP.

There are two 100Mb LANs. Machines are connected to both LANs. The load balancers are connected to the client LAN via separate connections. Client channel applications reside on the client's external LAN and have connections to the Internet, IVRs, and so on as they require.

Further Reading

- Performance of the CommSecure Markets Information System, *Dr James Mathew Farrow, CommSecure Limited, January 2002*

This White Paper is the property of CommSecure Limited (the "Company") and the information contained herein is protected by copyright. No part of this white paper may be reproduced in any form including photocopying, electronically or translation into a foreign language without the express written permission of the Company. Such unauthorised usage is a violation of Australian and International copyright law.

The Company considers the information contained in this White Paper to be reliable but it is not warranted by the Company. The document may also contain opinions of the Company, which are subject to change from time to time.
