

Strategic Systems Solutions

Project Profiles

Client Projects and Case Studies





Global Settlement System

Project Type: Production Support

Industry: Investment Banking

Technologies: Java, Oracle, Websphere

Implementation Type: Tier 2

The Situation

Our Financial Services client had staffed the Production Support group for its new settlement processing system with a team who had no prior IT experience. The group was comprised of employees who had Operations backgrounds only. Our client knew that these people were good at their current job, but soon learned that a good Operations employee does not necessarily make a good Production Support analyst. Within six months, the original DB team was replaced with a team of 8 SSS employees. The key problems that this SSS team faced were that there was no support model and the application was extremely unstable.

Challenge – No support model

The original Production Support team faced several issues including being fragmented across offices, not enforcing a central mechanism for reporting problems, and not maintaining a clear inventory of the systems that the team supported. In an attempt to provide global coverage, our client had implemented a ‘follow the sun’ model. Team members were located in various offices around the globe. When problems arose with the application, users could call a central number, but preferred instead to call support personnel directly or to send emails. The client’s Issue Tracking system was used as an afterthought, at best.

SSS implemented a 24 X 7 production support coverage model that is delivered from one location, the SSS Newcastle office in England. The team is led by a seasoned Production Support manager and is staffed with developers and DBAs. Clear communication was provided to the various user groups that problems would only be addressed if they were submitted through the central call in number or through the Issue Tracking application. Key contacts were also set-up for each business line where status could be communicated on urgent production problems.

Challenge – Unstable application

The Production Support team soon discovered that application code was routinely migrated into production with expected manual workarounds or without being properly tested. The client deemed time consuming nightly manual workarounds as a necessary evil to keeping pace with aggressive roll-out schedules.

After the infrastructure and procedures supporting the SSS Production Support team were in place, the team began tackling problems in the application itself. Production jobs that routinely failed or required extensive manual support were addressed first. The team also developed a set of required documentation and associated steps that developers must complete before promoting code to production. In the current environment, the Production Support team has an in-depth understanding of the Production environment and can offer suggestions to application developers and upstream systems feeding the application.

Current State

The team has grown to 18 people with 4 focusing on issue administration, 9 application specialists performing maintenance support, and 4 DBAs providing database support. The group now provides Management Information reports to the project’s management team and to the business users.



General System Description

- Global settlement system – European, US, and Asian implementations
- Web browser based - work flow / event driven philosophy
- 24x7x365 operation – high availability, STP
- \$100million annual budget
- 380 staff across UK, US, Asia
- 100+ SSS staff with 18 on this Production Support team.

Processing Stats

- 1,400+ daily batch jobs. The weekly process kicks off on Sunday at 10pm BST and run through midday on Saturday afternoon.
- 450+ intra day processes.
- ~300,000 US trades processed per night.
- 1,500 users (multi-language) based in London, NYC, Sydney, Tokyo, Hong Kong.
- 6-10 code production migrations each night, major Go Lives every 6-8 weeks.
- Weekend support includes routine Production code migrations and Disaster Recovery dress rehearsals.

Supported Application Details

- Web Technologies: XML, EJB, Servlets, JSP, Javascript, ASP, HTML
- Databases: Oracle, Sybase, SQL Server,
- Platforms: Websphere, Weblogic, Unix
- Open Source: JBoss, Struts
- Other: C++, Vitria, FTI, Mercator, Corba

Support Tools

- Batch Scheduler: Autosys
- Issue Escalation/Contact: Client Issue Tracker Application, Client Internal Hotline Number, Beepers?
- Database Access:
- Client Network Access:



Market Data Feeds

Project Type: Production Support

Industry: Investment Banking

Technologies: Perl, C++, Unix

Implementation Type: Tier 2

The Situation

Our Financial Services client wanted to move the production support of several vendor-provided market data feeds to a co-sourcing partner. SSS was chosen to provide this service and a team of four people was set-up in our Newcastle, England office. It was soon evident that the market data processing stream needed to be modified to improve overall controls both within the processing stream and at a project support level. As the support requirements grew, the team was also charged with providing a wide range of functional area and time coverage support, while maintaining a fairly small team.

Challenge – Minimal controls and inefficient processing stream

The client's processing stream was managed through cron jobs and there were few means to get an overall sense of how the nightly stream was progressing. Also, only a few jobs would provide a notification to the support team if it failed to complete successfully. To compound these problems, many jobs had significant performance issues. Finally, most programs had no data integrity checks and would allow invalid data to be loaded into the market data tables.

The SSS team implemented the batch scheduler AutoSys, which dramatically improved the group's control over the processing stream. When jobs failed, beeper notifications were sent to the appropriate team member on support. And finally, various web tools were developed to ensure that data was cleansed properly.

Challenge – No support model

The client had a complex set of market data requirements that needed to be supported in a 24 X 7 manner. The team's responsibilities included batch support and general job development/maintenance. The client wanted a cost-efficient means of providing this support while still providing high quality service to the user community.

The SSS team organized the Market Data batch jobs into five 'specialist' functional areas. Each functional area has two to four systems 'specialists' who perform the detailed job development/maintenance. Finally, the group is divided into four 'support groups' of three people each. Each support group has a mix of functional specialists. When a problem is raised, it is expected that a 'non-specialist' will be able to support the job stream if an appropriate specialist is not in the office. With this structure, the team can keep the number of development staff to a minimum (typically two per functional area) while still providing 24 X 7 production support coverage.

Challenge – Minimal coordination between developers

The client had no single point of contact for the Market Data feeds. Individual developers supported their own specific code and rarely worked with other developers to implement a consistent solution to production problems.

The Market Data team immediately began using the client's Issue Tracking system as the main repository of batch problems. Then as issues are raised across groups, they are discussed as a team to determine if there is a common source problem. For example, if there was an ftp firewall issue, the team will fix the root problem and prevent a more widespread issue from developing.



Current State

After the first set of feeds were rewritten and stabilized, additional feeds and systems have been added to the team's support responsibilities. Currently, the team has 12 people which provide 24 X 7 support for the various categories of market data. The group provides global coverage and interacts with support groups in the Far East, Europe, and US (east and west coast coverage).

General System Description

Support of daily loads of the following market data from external vendors:

- GIM - 723 traded indices comprising 140,122 components
- MSCI - 559 traded indices comprising 91,017 components
- SectorWatch - 424,213 sector products from 11 industry sources
- Corporate Actions - 8 vendors with approximately 446,000 notifications
- Reuters Real Time Product App - referential data for approximately 1 million products
- Options Clearing Corp – approximately 150,000 options with spikes on expiration weekends
- Bloomberg - 1 million Fixed Income securities and 293,000 Equity securities
- Fitch Ratings - 515,000 securities
- S&P Ratings - 135,000 securities
- Moodys - 197,000 securities
- FT Interactive - indicative data for 1.5 million municipal bonds
- Trepp - indicative data for 222,000 municipal bonds
- DTC - indicative data for 2.8 million securities
- Logical Information Machines (LIM) XMIM Loader - 3,000 customized products (e.g., oil and gas prices) and calculates firm-wide ValueAtRisk.

General Processing Stats

- GIM Loader runs six times a day and can run on an ad-hoc basis.
- MSCI runs three times a day (Far East, Europe, North America).
- All other loaders run once daily.

Supported Application Details

- Languages: ksh scripts, Perl, CGI Perl, C, C++
- Database: Sybase
- Platforms: SUN Unix

Support Tools

- Batch Scheduler: Autosys and Netcool.
- Issue Escalation/Contact: Client Issue Tracker Application, Client internal and external Hotline numbers, Beepers, Mobile phone. Staff on call and contact numbers are published on the Support Intranet page.
- Database Access: through standard Sybase tools, e.g. fsql, isql, wisql, PL/SQL. The team has also provided Web based tools to allow users "controlled" access to the DB.
- Client Network Access: via MS LAN in the SSS office then staff laptops are RAS/VPN enabled for home usage.



General Ledger Integration

Project Type: Post Merger Integration

Industry: Investment Banking

Technologies: Natural, Adabas, Mainframe

Implementation Type: Tier 1 & 2

The Situation

Our Financial Services client acquired an international Investment Bank and needed to integrate the acquired firm's General Ledger processing stream into its own. The project was complex due to the large number of international offices and numerous different systems that fed the Investment Bank's General Ledger. The project was also high profile, since our client announced an implementation date to the NYSE and the press for when it planned to merge the books of the two firms. A team of 8 SSS personnel, a project lead and 7 analyst/programmers, were selected to support this effort. The team worked alongside a team of 5 at the Investment Bank.

Challenge – Complex environment

The project scope included four international locations (London, Paris, Johannesburg and UK Subsidiaries) and 14 feeder systems, each using different GL interface templates. The team developed over 200 feeds to support the merger. The team worked with a new set of users who had different business rules, technologies and working practices than the 'parent' investment bank. The technologies, such as Unix, Windows, AS/400 and Mainframe, added further complexity with regard to data transfer and data conversion.

The SSS team first wrote detailed requirements based on the high level requirements received from the client. This process raised numerous issues which were resolved through user meetings, emails, and conference calls. Based on the final requirements, an Adabas-Natural mapping system was written, which runs off of user-maintained rules. The system takes an unlimited number of pre-defined inbound formats and maps them to the required G/L outbound format.

SSS developed several audit and exception reports and downloads to assist in the day to day running of the system. These reports highlight any data items that could not be mapped and have therefore been defaulted by the system. SSS also developed automated rule uploads written in Visual Basic. This tool allowed users to maintain rules locally on a PC and then upload them to the mainframe. This process saved our client significant time, as the PC allowed flexible maintenance and analysis before the final upload to the mainframe.

Another complexity in this environment was the inclusion of the Turkish Lira (TRL). This currency is infamous for its high exchange rate and often many systems cannot cope with the large number of significant digits. Some systems get around this issue by holding the amount against an alternative currency such as USD. The system was developed to support receiving large TRL values in various formats and mapping them accordingly.

Challenge - Aggressive, fixed implementation date

As stated earlier, the Implementation date for this project was seen as a market event. If this date was delayed there would have been significant negative press published about our client. To ensure that the Merger team met its deadlines, the team changed its traditional working style.

SSS implemented a shift pattern covering 18 hours per day. Many team members often worked longer hours to achieve the project goals. The team worked on site at the Investment Bank and were also on-call via pager/phone. Through this effort, the Merger team was able to successfully participate in three



User Acceptance Tests, three Implementation Dress rehearsals, and was ready for the ultimate Implementation date.

Current State

The implementation went smoothly on the Go Live date that was set one year earlier. The system is still running daily in Production across 4 locations and with over 220 feeds in place. The solution was so flexible that it is now being used as the mapping software for other new projects.

General System Description

The SSS team developed mapping software that allowed the acquired firm's feeder systems to interface with our client's General Ledger. The team provided file capture, reference tables, and mapping logic to deliver the target system files for multiple locations. The software managed and controlled the associated file transfers between platforms.

Key functionality of the application included the following:

- Batch Update of Mapping Tables
- Conversion Extracts
- Formatting/Reformatting, i.e. Unpacked/Packed
- Wrapping of Data, i.e. Dates
- Mapping of Data, i.e. GL Account Numbers
- Calculation/Recalculation of Amounts, i.e. TRL values, USD Equivalents
- Reporting: Maintenance, Exceptions, Reconciliation, etc.
- Online Maintenance and Enquiry of Mapping Tables and additional data
- Creation of Records & Reversing Records due to differences in accounting policies
- Splitting/Combining Records, i.e. Contact level to constituent parts - Principal Accrued Interest, etc. or vice versa

Processing Stats

- 200+ feeder systems
- Daily records processed (postings)
- Initial conversion of GL (all locations) : ~400,000 postings
- Post Conversion, average night (all locations, all feeder systems) : ~1,000,000 postings

Application Details

- Languages: Natural & COBOL
- Database: Adabas and QSAM
- Platforms: Mainframe and PC (rule maintenance)