

Tel Aviv Stock Exchange Case Study Disaster Tolerant Environment

The Tel Aviv Stock Exchange has taken an approach to disaster tolerance that balances availability, affordability, and data integrity. By periodically reviewing its disaster recovery plans, and regularly testing and updating its systems, TASE has made certain that it always had an appropriate level of protection.

Now, as its plans to expand the products and operations of its business, TASE has evolved its systems to deliver increased levels of failover and redundancy, with significantly reduced complexity.

Section 1 Business and Technology Overview

TASE Profile

The Tel Aviv Stock Exchange (TASE), established in 1953, offers international investors the opportunity to participate in Israel's growing, technology-driven economy. In recent years, many distinguished investment houses have become active and influential in the Israeli stock market.

TASE handles trading of Israeli securities and derivatives, some of which are also traded on NYSE and Nasdaq.

To meet the needs of the international investment community, the TASE has responded with a number of technological improvements, putting itself on an equal footing with the world's most advanced stock exchanges. In 1997, the TASE introduced the fully automated Tel Aviv Continuous Trading (TACT) system, an integrated platform that now supports the trading of all listed securities and derivative products.

The TASE enables member firms to offer on-line trading services to investors through the Internet. The new technologies helped boost trading volume and reduce volatility, thereby mitigating the costs and the risks of trading.

The TASE's high regulatory and technological standards have been recognized by the U.S. Securities and Exchange Commission, which has named the TASE a designated offshore securities market. The TASE is an active member of both the International Federation of Stock Exchanges (FIBV) and the International Options Markets Association (IOMA).

TASE Infrastructure

Servers

To support its securities and derivatives trading operations, the TASE uses HP AlphaServer systems running OpenVMS clusters. A cluster of two AlphaServer 4000 systems run the securities trading application and a similar cluster of two AlphaServer 4100 systems run the derivatives trading application.

These OpenVMS clustered systems match the bid and ask prices received from the trading members, notifying them whether the trade has been accomplished, how many shares were bought or sold, and at what price. At the same time, the system publishes the price and volume of the last trades, along with the new bid and ask prices, to the members' dealing rooms and to data vendors' systems.

Each cluster consists of two servers: a primary and a hot backup. Both servers are active during the day, but only the primary is connected to the members' systems and dealing rooms. The primary server duplicates each transaction and sends it to the hot backup. Therefore, the hot backup contains a full copy of the trading activity. In case the primary server fails, the hot backup can take its place within a few minutes.

To communicate with its members, TASE uses AlphaServer DS20E systems running Tru64 UNIX and System Network Architecture (SNA) software and TCP/IP. Two AlphaServer DS20E systems are allotted for securities trading and two for derivatives trading, in each case with one system active and the other reserved for backup. In the event that the production Tru64 UNIX system fails, TASE boots the second system, which then mounts the disks and assumes the identity of the failed system.

The SNA software also enables the front-end AlphaServer DS20E systems to communicate with TASE's IBM mainframe. At the end of each day, trading results are passed to TASE back office systems where the clearing takes place. At the beginning of the next day, the IBM mainframe downloads the most recent closing prices—reflecting the latest trading in other markets-- of all listed stocks to the OpenVMS systems.

Network

The AlphaServer 4000/4100 and DS20E servers are connected via LAN using TCP/IP. TASE does not consider the Internet secure enough to connect to the trading systems.

Communications with members is over a WAN supplied by Bezeq, the telephone company.

In the near future TASE is going to introduce a WIRELESS WAN as a backup channel on the same infrastructure.

The Tru64 UNIX system communicates with the members over SNA or TCP/IP, depending on the member's needs. It accepts trade orders, passing them to the trading queue on the OpenVMS system. There, the trades are resolved against other trade orders, and the results forwarded to the Tru64 UNIX system, which notifies each member.

Storage

TASE uses SCSI adapters on the AlphaServer systems connected to StorageWorks disk enclosures, so the disks are simple SCSI devices.

The two OpenVMS clustered servers are connected to local shared SCSI storage. The Tru64 UNIX systems are not clustered, but they are connected to a shared SCSI bus.

To address their DT requirements TASE opted for a log shipping environment in which journal logs are copied and stored separately at preset intervals. This an economical environment that provides reasonable RPO and RTO based on the current demands of the business model.

The primary trading computer logs every transaction to RAM. In the event of a disaster that results in system loss (both primary and hot backup), TASE recreates the state of the market using the backed up journal logs. The system starts from the market's opening state, sequentially using the backup logs to determine the latest pricing. At that point, trading can resume.

Section 2

Business Continuity Approach

Existing Business Continuity Plan

The TASE's business continuity plan addresses both local server failure as well as more extensive disaster. If the Tel Aviv site experiences a system failure, there would be a pause in trading, but no significant effect on the market.

Through a carefully supervised process, the TASE IT staff manually fails over to its co-located backup system. To quickly restore trading, the backup systems recreate the market's current state by reading the log files that were stored to disk on the backup OpenVMS platform. The entire Recovery Time Objective (RTO) is up to ten minutes, down from a recent standard of twenty minutes.

In the event of a failure of the entire Tel Aviv computer room or building, recovery would take longer, anywhere from 1-5 hours. In this case, the state of the market would be lost, since it is stored in memory and on disk only at Tel Aviv. Thus a major disaster could result in RPO and RTO that, while they had been acceptable in the past, may not serve them well in the future.

To address this level of disaster, the TASE maintains cold standby backup servers, storage and networking equipment at an HP facility in Raanana. This equipment is also used for testing software upgrades and quality assurance.

In case of disaster, depending on the state of communications in the Tel Aviv command center, TASE could either deploy the backup systems with the help of HP staff at Raanana, or relocate their operations to the Raanana site.

The Raanana Disaster Recovery site enables the TASE to provide full service to their members, although without the same level of redundancy that they have at Tel Aviv. The DR site has only one OpenVMS, UNIX and NT server per application, with no backup servers. The Raanana site also has office space, documentation, and telephone lines, enabling a minimal staff to support on-going activities.

At Raanana, the DR plan includes restoring the software from the backup tapes, reconfiguring the network, and notifying the members before trading resumes. In the event of a disaster of this magnitude, TASE's RTO would be the morning of the next business day.

TASE's Recovery Point Objective (RPO) is full recovery to the last trade logged before server or network failure. In the event of disaster, RPO is to the beginning of the trading day, adjusted for whatever happens in out-of-hours and overseas trading.

A senior member of the TASE's IT department is responsible for its Disaster Recovery plan, testing and updating it as necessary. The mainframe, OpenVMS and Tru64 UNIX system managers are also members of the DR team.

Evolving the Disaster Recovery Plan: The TASE vision

TASE has continuously evolved its disaster recovery capabilities to meet changing needs. For example, in the event of a system failure, the OpenVMS production server could failover only once a day to the backup server. Recently, TASE modified its applications to enable the backup server to assume the primary system role as often as necessary during a trading session.

As a next step in evolving its disaster recovery plan, the TASE anticipates designing automatic failover capability into its systems. To implement this technology, the stock exchange will upgrade its software, especially to eliminate the vulnerability of investors who take positions in the derivatives market.

The TASE is also considering moving from a logging system to conduct trades to a database system that would reside on disks. A database system would have to be fast enough to support the rate of the TASE's trading activity, a concern of the stock exchange. However, database technology offers the advantage of more quickly determining the market's current state in the event of system failure or disaster.

For the past five years, the TASE has evolved its systems with the goal of achieving high availability for its markets. Next, the stock exchange would like to move to a storage replication system that uses disks to manage the market, with improved recovery point and recovery time capabilities. By storing the data in a shared disk environment so that all the servers could access it, recovery would be easier and faster.

“In the long run, we’re looking at storage technology to give us high availability and disaster-tolerance,” says Jonathan Shemesh, senior vice-president of computers and operations at the TASE. “Storage will be one of the most important elements to enable us to achieve the recovery time standards that we’ve set.”

“The system has to be capable to deal with, firstly, application failure, secondly, any server failure... and thirdly, any storage failures,” says Mr. Shemesh. “HP’s hardware engineers have developed an effective disaster-tolerant storage solution, but the applications we use will require enhancements before they can be considered fault tolerant.”

The TASE also hopes to mirror the disk data in real-time to the Raanana site, perhaps using StorageWorks Data Replication Manager, so that, even in the event of disaster, trading could continue without loss of data.

HP services and pre-sales systems experts having been working with the TASE to evolve their business continuity and DT model. The planning and analysis they have already done will help TASE evolve in stages, while retaining compatibility and investment protection.

Any major technological changes at TASE must be approved by its members, who comprise its board of directors.

Business Critical Consulting & High Availability Partnership

TASE has a vision of extending their trading hours and services, which would place new limits on its disaster recovery window. Capabilities such as automatic failover in case of server failure or disaster would become increasingly necessary. HP’s business critical consultants have been working with the stock exchange to develop a schedule that meets its needs to evolve its business continuity technology.

HP has been working closely with the TASE, providing support with change management, configuration management, patch and blitz notification, priority escalation routes, as well as providing a dedicated support team and monthly status meetings with HP service managers.

As a result, the TASE recently upgraded from AlphaServer 1200 series platforms to the AlphaServer DS20E systems for trading and implemented a number of enhancements to improve system stability.

HP’s business critical team has completed an availability review, performed an Internet Security Healthcheck, and proposed a business critical high-availability partnership.

Currently, the TASE is evaluating an HP proposal to deliver significantly faster recovery time in the event of disaster. The first stage that they are considering is a single-site SAN for the Tel Aviv location using Fiber Channel controllers. The next stage would be extending this SAN to make it multi-site.

“The future of the TASE will include longer trading hours and new instruments to be traded,” says Mr. Shemesh. “We’ll need fault-tolerance and high-availability. I think that with a combination of using smart applications, a model server configuration, and adding advanced storage, we can achieve it.”

###

Compaq Computer Corporation is a wholly owned subsidiary of Hewlett-Packard. IBM is a trademark of International Business Machines, Inc. UNIX is a registered trademark of The Open Group in the US and other countries.

© 2002 Hewlett-Packard Company. All Rights Reserved. All other names and products are the property of their respective companies. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

October 2002