

Highly Available Systems: Concepts and Practice Two-Day Seminar

Seminar Description

This one-day seminar describes how highly available systems with annual downtimes measured in minutes or seconds can be implemented. The secret to high availability is fast failover to a redundant backup system.

An introduction to basic availability concepts leads to methods for implementing highly available systems. Fundamental to all highly available systems is data replication to a remote site. Various data replication techniques are studied, and commercially available data replication engines are described.

Highly available architectures that support a wide range of availability characteristics are presented. They include active/active systems that can provide uptimes measured in centuries by recovering from a fault in seconds.

Several case studies demonstrate successful implementations. Sprinkled throughout the seminar are many “never again” educational stories of disastrous failures of critical systems.

We will tailor the content of the seminar to meet your needs and interests.

Seminar Objectives

Attendees can expect to learn the following:

- Real examples of what can go wrong.
- Basic availability concepts
- The impact of redundancy on availability

- The use of data replication for database synchronization
- Data replication products
- How RPO and RTO is affected by the choice of data-replication technique
- Highly available architectures
- Active/active systems
- Eliminating planned downtime
- Real-life examples of highly available systems in practice

Prerequisites

Familiarity with transaction-processing systems and knowledge of simple algebra is recommended.

Instructor Biography

Dr. Bill Highleyman brings more than 40 years' experience in the design and implementation of mission-critical computer systems to his position as Chairman of The Sombers Group. Dr. Highleyman, a graduate of Rensselaer Polytechnic Institute and MIT, earned his doctorate degree in electrical engineering from Polytechnic University. He has published extensively on availability, performance, testing, and middleware issues. He is the author of “Performance Analysis of Transaction Processing Systems,” published by Prentice-Hall, and is coauthor of the three-volume series, “Breaking the Availability Barrier.” Dr. Highleyman holds several patents and publishes the free monthly Availability Digest available at www.availabilitydigest.com.

Highly Available Systems: Concepts and Practice Day 1

9:00 – 10:30: Concepts in Availability

Continuous availability versus high availability
MTR, MTBF, Availability, and their relationship
The 9s game – Measuring availability as 9s
Recovery Time Objective (RTO) and Recovery Point Objective (RPO)
The impact of redundancy on availability
Calculating availability
Node failures
Effects of failover times and faults
Active/active systems – an introduction
Comparison to clusters

10:30 – 10:45: Break

10:45 – 12:15: Data Replication

Synchronizing database copies with data replication
Asynchronous replication:
Replication latency
Data loss following a node failure
Data-collision detection and resolution
Synchronous replication:
Application latency
Network deadlocks
The importance of redundant, reliable networks
Replication does not eliminate backups

12:15 – 1:15: Lunch

1:15 – 2:45: Data Replication Engines

The replication engine check list
Asynchronous replication engines
Shadowbase
Goldengate
IBM Global Mirror
Sybase, SQL Server, MySQL
Synchronous replication engines
Shadowbase Plus SR
IBM Parallel Sysplex
IBM Metro Mirror
SAN Replication
Grow your own

2:45 – 3:00: Break

3:00 – 4:30: Highly Available Architectures

Fault-tolerant systems (single node)
HP NonStop
Stratus
Active/standby systems:
Magnetic tape
Virtual tape
Active/passive systems:
Unidirectional async replication
Unidirectional synchronous replication
Active/active systems:
Bidirectional async replication
Bidirectional synchronous replication
Eliminating planned downtime
Fast failover
User redirection
Transaction redirection
Other advantages of active/active systems

Highly Available Systems: Concepts and Practice Day 2

9:00 – 10:30 Cloud Computing

What is the cloud?
Advantages of the cloud
Public clouds
Private clouds
Hybrid clouds
Notable cloud failures
 Amazon Web Services
 Google Apps
 Windows Azure
 Salesforce.com
Cloud SLAs
Connecting to the cloud
 Internet reliability
The compute utility
The appropriate role of the cloud

10:30 – 10:45: Break

10:45 – 12:15 Cyber Attacks

Types of cyber attacks
 Spear phishing
 Trojans
 Worms
 Viruses
 Key loggers
 Spyware
 Ransomware
Mobile device threats
Distributed Denial of Service (DDoS)
 The biggest threat to availability
 Botnets
 Types of DDoS attacks
 Recent examples
 Mitigation

12:15 – 1:15: Lunch

1:15 – 2:45 Never Again

Examples of massive system failures
 Human error
 Failed upgrades
 Failover faults
 Internet
 Intranet
 Environmental
Lessons learned

2:45 – 3:00: Break

3:00 – 4:30 Active/Active Systems in Action

Case studies
 Financial
 Communications
 Control systems
 Other case studies