

## Highly Available Systems: Concepts and Practice Three-Day Seminar

### Seminar Description

This three-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 copy synchronization

- Data replication products
- How RPO and RTO is affected by the choice of data-replication technique
- Reliable networks
- Highly available architectures
- Active/active systems
- Eliminating planned downtime
- The impact of virtualization
- Cloud computing
- 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 monthly free Availability Digest ([www.availabilitydigest.com](http://www.availabilitydigest.com)).

## Highly Available Systems: Concepts and Practice Day 1

### **Part 1 – Concepts in Availability**

Continuous availability versus high availability  
Active/active systems - an introduction  
Availability, failure rates, and reliability  
The basic availability equation  
The 9s game  
Basic availability relations  
    Serial systems  
    Parallel systems  
    Reliability block diagram  
Availability of dual-node systems  
    Dual-node failure  
    Failover time  
    Failover faults  
Dual-node comparisons  
    Clusters  
    Active/active systems

### **Part 2 – Multinode Systems**

The three Rs in MTR  
    Repair  
    Recovery  
    Restore  
Repair strategies  
Multinode systems  
    Process allocation  
    Sparing  
Effect of hardware/software failures  
Approximation errors  
Case study

### **Part 3 – Active/Active Systems**

Availability – a review  
    Failures  
    Failovers  
    Failover faults  
RPO and RTO  
Active/passive systems  
Active/active systems  
Why are active/active systems reliable?  
    Redundancy  
    Isolation  
    Dispersion  
    Failover  
Active/active topologies  
    System splitting  
    Disk farms  
    Symmetric systems  
    Asymmetric systems  
    Partitioned systems  
Active/active networking  
    Redundancy  
    Split-brain mode  
    Communication SLAs  
Fast failover  
    User redirection  
    Router redirection  
    Sever redirection  
    DNS redirection  
Application issues  
    Impediments  
    “Sizzling-hot” standby  
Heterogeneous nodes  
Case study

## Highly Available Systems: Concepts and Practice Day 2

### **Part 4 – Database Replication**

- Purpose of database replication
- Required characteristics
  - Bidirectional
  - Ping-ponging
  - Capacity expansion
  - Distributed management
  - Automatic failure recovery
  - Create, compare, synchronize
- Asynchronous replication:
  - Replication latency
  - Data loss
  - Data collisions
    - Detection
    - Resolution
  - Calculation of data collision rates
- Synchronous replication
  - Application latency
  - Network transactions
  - Coordinated commits
  - Recovery
- Transaction replication
- Hardware replication
- Replication network
  - Split-brain mode
  - Communication SLA
- Online copy
- Validation and verification
- Database of record

### **Part 5 – Data Replication Engines**

- The replication engine check list
- Asynchronous replication engines
  - Shadowbase
  - Goldengate
  - DRNet
  - Oracle Streams
  - IBM Global Mirror
  - Tungsten Replicator
  - Double-Take
  - Sybase, SQL Server, MySQL
- Synchronous replication engines
  - Shadowbase Plus SR
  - OpenVMS split-site clusters
  - IBM Parallel Sysplex
  - IBM Metro Mirror
- Grow your own

### **Part 6 – Other Advantages of Active/Active Systems:**

- Disaster tolerance for free
- Fast and reliable failover
- Elimination of planned downtime
- Data locality
- Application scaling via symmetric expansion
- Application scaling via asymmetric expansion
- Load balancing
- Lights-out operation
- Efficient use of all capacity
- Elimination of failover decision time
- Risk-free failover testing

## Highly Available Systems: Concepts and Practice Day 3

### **Part 7 – Comparison to Clusters**

- Cluster architectures
  - Cluster resource group
  - Active/standby clusters
  - “Active/active” clusters
- Cluster failover
- Zero-downtime upgrades
- Split-brain mode
- Application constraints
- Application scaling
  - Oracle RAC
- Cluster availability
- Disaster tolerance
- Review of active/active
- Comparison to active/active
  - Availability
  - Split brain
  - Data sharing
  - Application scaling
  - Disaster recovery
  - Heterogeneity
  - Zero downtime upgrades
  - Maturity

### **Part 8 – Cloud Computing**

- What is the cloud?
- Advantages of the cloud
- Notable cloud failures
  - Compute cloud
  - Storage cloud
  - Hosting cloud
  - Software-as-a-service
  - Platform-as-a-service
- Private clouds
- Cloud SLAs
- Connecting to the cloud
  - Internet reliability
- The compute utility
- The appropriate role of the cloud

### **Part 9 – Active/Active Systems in Action**

- Financial
  - Bank-Verlag
  - BankServ
  - Fifth Third Bank
  - Banco de Credito
  - Handelsbanken
- Communications
  - Telecom Italia
  - HP OpenCall INS
  - HP Home Location Register
  - Vodacom
- Control systems
  - QEI TDMS-PLUS
  - Chicago MTA
- Other case studies
  - Internet service provider
  - European blood bank
  - New York Racing Association