

## **Active/Active Systems: Theory and Practice One-Day Seminar**

### **Seminar Description**

This one-day seminar describes how active/active systems comprising two or more servers can provide uptimes measured in centuries. An active/active system is a network of independent, geographically-distributed processing nodes cooperating in a common application. Each node has access to an up-to-date copy of the application database. Should a node or a database copy fail, all that needs to be done is to switch over that node's users to a surviving node. Recovery is in subseconds to seconds.

The underlying concept of active/active systems is "let it fail, but fix it fast." If users never notice an outage, then in effect, an outage hasn't occurred.

### **Seminar Objectives**

Attendees can expect to learn the following:

- Basic availability theory
- The impact of redundancy on availability
- The architecture of active/active systems
- The use of data replication to keep database copies synchronized
- How RPO and RTO is affected by the choice of data-synchronization technique
- Data replication products
- Eliminating planned downtime

- Other advantages of active/active architectures
- How active/active systems compare to clusters
- Real-life examples of active/active systems in practice

### **Prerequisites**

Familiarity with transaction-processing systems is required. A knowledge of simple algebra and basic probability theory is recommended though not necessary.

### **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 Availability Digest, which focuses on topics related to continuous availability. The Availability Digest is available at [www.availabilitydigest.com](http://www.availabilitydigest.com).

## **Active/Active Systems: Theory and Practice Seminar Outline**

### **9:00 – 10:30: Concepts in Availability**

MTR, MTBF, Availability, and their relationship

The 9s game – Measuring availability as 9s

The impact of redundancy on availability

Failure modes – what are they and what is their affect on availability

Calculating availability

### **10:30 – 10:45: Break**

### **10:45 – 12:15: Active/Active Systems**

What is an active/active system?

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

Recovery Time Objective (RTO) and

Recovery Point Objective (RPO)

The importance of redundant, reliable networks

Other uses of data replication:

- Disaster recovery

- Heterogeneous application integration

Active/active systems versus clusters

### **12:15 – 1:15: Lunch**

### **1:15 – 2:45: Other Advantages of Active/Active Architectures**

- Elimination of failover decision time

- Disaster tolerance for free

- Improved data locality

- Eliminate planned downtime

- Efficient use of all capacity

- Risk-free failover testing

- Application scaling via symmetric expansion

- Application scaling via asymmetric expansion

- Load balancing

- Lights-out operation

### **2:45 – 3:00: Break**

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

Replication products:

- Oracle Streams

- GoldenGate

- Shadowbase

- DRNet

Financial case studies

Telecommunication case studies

Other case studies