

Highly Available Systems: Concepts and Practice Three-Day Seminar

Seminar Description

This three-day seminar describes how redundant computing systems comprising two or more servers can provide uptimes measured in years.

An important class of highly available systems that is described is active/active systems, which are networks 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 highly available 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
- The use of data replication
- RPO and RTO
- Fast failover is the key
- Active/active systems
- Reliable networks
- Achieving fast failover
- Eliminating planned downtime
- Other advantages of active/active architectures
- Data replication products

- How active/active systems compare to clusters
- The threat of cyberattacks
- Distributed Denial of Service (DDoS) attacks
- Cloud computing
- Real-life horror stories
- Real-life examples of active/active systems in practice

Prerequisites

Familiarity with transaction-processing systems is helpful. 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 N.Y. 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.

**Active/Active Systems: Theory and Practice
Day 1**

Part 1 – Concepts in Availability

Continuous availability versus high availability
Availability, failure rates, and reliability
RPO and RTO
The basic availability equations
The 9s game
Basic availability relations
Serial systems
Parallel systems
Reliability block diagram
Active/active systems - an introduction

Part 2 – Multinode Systems

Availability of dual-node systems
Dual-node failure
Failover time
Failover faults
High availability is fast failover
Dual-node comparisons
Clusters
Active/active systems
The three Rs in MTR
Repair
Recovery
Restore
Repair strategies
Sparing
Approximation errors
Case study

Part 3 – Active/Active Systems

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

Part 4 - Fast failover

Active/active networking
Redundancy
Split-brain mode
Communication SLAs
Methods for fast failover
User redirection
Router redirection
Server redirection
DNS redirection
Application issues
Impediments
“Sizzling-hot” standby
Heterogeneous nodes
Case study

Active/Active Systems: Theory and Practice Day 2

Part 5 – 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
Synchronous replication
 Application latency
 Recovery
Transaction replication
Hardware replication
Replication network
 Split-brain mode
 Communication SLA
Online copy
Validation and verification
Database of record

Part 6 – 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

Part 7 – 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

Part 8 – 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

**Active/Active Systems: Theory and Practice
Day 3**

Part 9 – 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

Part 10 – 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
 Recent examples
 Botnets
 Types of DDoS attacks
 Mitigation

Part 11- Never Again

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

Part 12 – Active/Active Systems in Action

Financial
Communications
Control systems
Other case studies