

Active/Active Systems: Theory and Practice Five-Day Seminar

Seminar Description

This five-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
- Fault-tolerant systems
- The architecture of active/active systems
- The use of data replication to keep database copies synchronized
- How RPO and RTO are affected by the choice of data-synchronization technique
- Reliable networks
- Data-replication products

- Eliminating planned downtime
- Other advantages of active/active architectures
- How active/active systems compare to clusters
- Virtualization
- Cloud computing
- Real-life horror stories
- 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 RPI 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," 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
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

**Active/Active Systems: Theory and Practice
Day 2**

Part 3 – Failure-State Diagrams

Definition of failure state diagrams
Simple example of a redundant system
Application to active/active systems
Application to repair strategies
Application to repair and restoration
Application to recovery
Application to hardware/software faults

Part 4 – Active/Active Systems

Availability – a review
Failures
Failover
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

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

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
 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 6 – 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

Active/Active Systems: Theory and Practice Day 4

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
 - Maturity

Part 9 – Fault-Tolerant Systems

- Active/active systems
 - HP OpenVMS Split-Site Clusters
 - IBM Parallel Sysplex
- Distributed systems
 - HP NonStop
- Others
 - Stratus ftServer
 - Stratus Avance
 - Marathon everRun

Active/Active Systems: Theory and Practice Day 5

Part 10 – Virtualization

- What is virtualization?
- The drivers for virtualization
- Virtualization architectures
 - Operating system virtualization
 - Bare metal virtualization
 - Paravirtualization
- Virtualization and availability
 - Virtual machine failover
 - Clustering
 - Server pooling
- Disaster recovery
- Virtualization availability products
 - VMotion from VMware
 - LiveMigration from Virtual Iron
 - Hyper-V from Microsoft
- The requirement for fault tolerance

Part 11 – 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 12 – Never Again

- Major bank down for weeks
- Hubble failover fault
- IRS goof costs taxpayers \$300m+
- London Stock Exchange
- PayPal upgrade – no fallback
- Amtrak train wreck
- Alaska's \$38 billion key stroke
- The flying cable
- Twitter's DDoS attack
- Microsoft's Sidekick loses all data

Part 13 – 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