HIGH AVAILABILITY (1 hour 50 minutes)

Roadmap for Creating an “Always On” Environment

Janet Weber
First Union National Bank
1100 Corporate Center Drive
Raleigh, NC 27607
(919)852-6821 Direct
(919)852-6830 Fax
janet.weber@firstunion.com

HIGH AVAILABILITY

Roadmap for
Creating an
“Always On”
Environment

HPWorld 2001

The Journey

- Definitions
- Objectives
- Reliability
- Redundancy
- Availability
- Policies/Procedures
- Disaster Recovery

Definitions

- Availability
  - Availability is the percentage of time the system is performing normal business.

$$\frac{\text{Total elapsed time} - \text{Total downtime}}{\text{Total elapsed time}} \times 100\%$$

  - Elapsed time is continuous and includes operating time and downtime.

Definitions (cont.)

- Downtime
  - Downtime is the duration of an outage of normal business, planned and unplanned.
  - Average downtime is a measure of the outage time per failure.

$$\text{Average Downtime} = \frac{\text{Total downtime}}{\text{Total number of outages}}$$

Definitions (cont.)

- Planned downtime
  - Backups
  - Software upgrades
  - Hardware/Firmware maintenance
  - Patching
  - Moves
  - Migrations
  - Configuration

Definitions (cont.)

- Unplanned downtime
  - Failures (hardware, software, network, power or environment)
  - Database/Application error
  - User error
  - Human error
  - Natural disaster

Definitions (cont.)

- Single points of failure (SPOF)
  - Points where normal business service can be broken
  - Points where standby, alternate or redundant components are not available
  - People can also be single points of failure

Definitions (cont.)

- Mean Time Between Failures (MTBF)
  - MTBF is used to predict reliability, based on past performance.

$$\text{MTBF} = \frac{\text{Total elapsed time}}{\text{Total number of failures}}$$

Definitions (cont.)

- Mean Time to Repair (MTTR)
  - MTTR is used to predict downtime, based on average repair time.

$$\text{MTTR} = \frac{\text{Total of all repair times}}{\text{Total number of failures}}$$

Definitions (cont.)

- High availability
  - High availability is the HIGH percentage of time the environment is performing normal business
  - High availability depends on your perspective
    - Desktop
    - Network
    - Server
    - Application

Definitions (cont.)

Availability	Uptime	Downtime
100%	8760 hrs/yr	0 hrs/yr
99.999%	8760 hrs/yr	5 min/yr
99.99%	8759 hrs/yr	1 hr/yr
99.95%	8755 hrs/yr	5 hrs/yr
99.9%	8751 hrs/yr	9 hrs/yr
99.5%	8716 hrs/yr	44 hrs/yr
99.0%	8672 hrs/yr	88 hrs/yr

Definitions (cont.)

- Fault tolerance
  - Faults are anticipated events, and thus tolerated
  - Fault tolerance is a combination of hardware and software implementations
  - In the event a failure occurs, a backup component or procedure can immediately take its place with no loss of service

Definitions (cont.)

- Disaster tolerance
  - Disaster tolerance is being able to recover quickly from a disaster
  - Disaster tolerance depends on different components
    - redundant hardware
    - data replication
    - geographic separation
    - partial or complete recovery automation
    - well-defined recovery procedures

Objectives

- Determine the requirements for availability
- Assess risks
- Determine acceptable risks
- Design the system to meet acceptable downtime requirements
- Reduce or eliminate single points of failure

Requirements

- What is the nature of the business?
- What is an outage?
- What does an outage cost?
- How long can the system be down for any given time? For the year?
- Should the design focus on minimizing the number of outages or minimizing the duration of the outages?

Requirements (cont.)

- How highly available?
- What are the components for the system?
- What components need to be redundant?
- What components can fail?
- What are the backup and recovery strategies? What do they protect against?
- What is the disaster recovery plan? What does it protect against?

Requirements (cont.)

- What is the total usable time of the system?
- What is the total planned downtime?

Assessing Risks

- Hardware reliability
  - HIGHEST Mean Time Between Failures (MTBF)
  - LOWEST Mean Time to Repair (MTTR)
  - Both ratings must be considered

Assessing Risks (cont.)

- Operating system reliability
  - System recovery capabilities
  - Deallocation
  - Dynamically loadable kernel modules
  - Dynamically tunable kernel parameters
  - HP-UX partitions

Assessing Risks (cont.)

- Operating system reliability (cont.)
  - Process Resource Manager (PRM)
  - Workload Manager (WLM)
  - Online Addition and Replacement (OLAR)
  - Instant Capacity on Demand (iCOD)
  - Call home capability

Assessing Risks (cont.)

- Redundant root disks
  - Protects operating system components
  - Use disk mirroring (MirrorDisk/UX)
  - Choose hot swappable disks
  - Implement Journaled File System (JFS)
    - JFS features on-line configuration changes and fast recovery from failures
    - Ability to grow filesystems on-line is important

Assessing Risks (cont.)

- Uninterruptible Power Supply
  - In power loss, can sustain power long enough to allow for graceful shutdown
  - In power loss, can sustain power long enough to survive the power outage
  - In power spikes and dips, can protect system from power transients
  - Cost is directly related to amount of power needed and the duration supplied

Assessing Risks (cont.)

- Redundant data disks
  - Disk mirroring
    - Software mirroring (MirrorDisk/UX) provides redundant copies of information which could be used for backup and recovery purposes with little or no interruption of service
  - Redundant Array of Inexpensive Disks (RAID)
    - Heavily cached RAID arrays provide a disk subsystem with nearly all-redundant components, hot failover and replacement, and virtually eliminate downtime due to failure and repair.

Assessing Risks (cont.)

- Redundant disk links
  - Dual path between the server and the disk subsystem provides protection against disk interface failure
  - IO traffic will dynamically switch to the redundant path

Assessing Risks (cont.)

- Storage Area Network (SAN)
  - SANs provide high availability, high performance, security, flexibility, scalability, and manageability
  - Storage capacity, tape and disk, is used as one large pool providing any-to-any connectivity
  - IO traffic is separated from existing network

Assessing Risks (cont.)

- Storage Area Network (SAN) (cont.)
  - HP Surestore XP arrays
    - SAN Manager Device Management (DM)
    - SAN Manager LUN Management (LM)
  - EMC Symmetrix
  - STK silos

Assessing Risks (cont.)

- Redundant networking
  - Networks are subject to congestion and other problems beyond the control of the SA
  - Redundant network components reduce the risk of network failure when they are configured to dynamically failover and load balance

Assessing Risks (cont.)

- Redundant networking (cont.)
  - Auto Port Aggregation (APA)
  - Cisco Catalyst Family switches/7200 Series routers
  - AT&T Ultravailable Solutions provides secure multi-site clustering, completely managed and monitored
  - Brocade SAN switches

Assessing Risks (cont.)

- System failover
  - Software feature which allows a specified application or workload to migrate from one server to another in case of a failure
  - MC/ServiceGuard
  - MC/ServiceGuard Manager
  - Metro/ContinentalClusters
  - ServiceControl Manager (SCM)
  - EMC Symmetrix Remote Data Facility (SRDF)

Assessing Risks (cont.)

- Application
  - Oracle Parallel Fail Safe
    - Oracle Parallel Server with MC/ServiceGuard OPS edition
    - Active/Standby cluster configuration
  - BEA Tuxedo

Assessing Risks (cont.)

- Application (cont.)
  - HP Openview
    - Network Node Manager
    - Omniback
    - VantagePoint
    - MeasureWare

Assessing Risks (cont.)

- Stratus Continuum
  - “The World’s Most Reliable Servers”
  - HP PA-RISC symmetric multiprocessing technology
  - Supports HP-UX operating System
  - Proven fault tolerant architecture with multiprocessing, fast onboard memory, and dedicated I/O processors
  - Guaranteed continuous availability with duplex, self-checking hardware and logic

Assessing Risks (cont.)

- SuperDome
  - Keystone to 5nines:5minutes program
  - Each partition is equivalent of a traditional standalone system
  - Each partition comes with core I/O, other I/O and LAN connections
  - Each partition connects to boot devices, data disks, removable media (DVD-ROM and/or DAT)

Assessing Risks (cont.)

- SuperDome (cont.)
  - Redundant components exist in each partition
    - Disk and LAN interfaces
    - Heartbeat LANs
    - Boot devices via mirroring
    - Critical data via mirroring
    - LAN protection

Assessing Risks (cont.)

- SuperDome (cont.)
  - Any partition that is protected with MC/ServiceGuard can be configured with
    - A standalone system
    - Another partition with the SuperDome cabinet
    - Another SuperDome
  - Any partition that is protected with MC/ServiceGuard contains as many redundant components as possible to further reduce the chance of failure

Assessing Risks (cont.)

- Policies and Procedures
  - Policy dictates management of the system, and must reflect business needs of the application supported
  - Procedures implement this policy and clearly define how to maintain availability
  - Can distinguish success from failure

Assessing Risks (cont.)

- Support contracts
  - Hardware response time
    - Priority Plus On-Site Support (24x7)
      - provides four-hour response, 24-hours per day, 7-days per week, including holidays.
    - Priority On-site Support (8x5)
      - provides four-hour response, Monday through Friday, from 8:00 am to 5:00 pm local time, excluding holidays.
    - Next Business Day
      - provides service on the next working day after the call is received, Monday through Friday, from 8:00 am to 5:00 pm local time, excluding holidays.

Assessing Risks (cont.)

- Support contracts (cont.)
  - Hardware repair commitment service
    - Hardware Call-to-Repair Commitment service provides HP's highest level of reactive 24x7 hardware support with a commitment to repair the customer's hardware within a maximum of six hours.

Assessing Risks (cont.)

- Support contracts (cont.)
  - Software support includes software telephone assistance and software updates. The customer may choose the type of telephone support
    - normal business hours
    - around the clock telephone support (24x7).

Assessing Risks (cont.)

- Support contracts (cont.)
  - Personalized Systems Support (PSS)
    - HP Personalized Systems Support (PSS) is an on-going, personalized technical relationship with HP focused on meeting the customer's priority needs to proactively maintain and extend their IT environment. An assigned Account Support Engineer (ASE) serves as your primary contact with the HP Support organization.

Assessing Risks (cont.)

- Support contracts (cont.)
  - Critical Systems Support (CSS)
    - HP Critical Systems Support (CSS) offers industry leading response and repair, and flexible and modular preventive services to fit your specific needs, including consulting to reduce problems and downtime, and increase effectiveness. An Account Support Engineer (ASE), who knows your business and system, leads a team of assigned experts in supporting high availability computing environments. Your ASE works with your team on technical and operational issues to help reduce the frequency of systems failures.

Assessing Risks (cont.)

- Support contracts (cont.)
  - Business Continuity Support (BCS)
    - Business Continuity Support for mission critical environments includes assigned teams, proven processes, and customized tools and ensures that no one has higher priority - or more visibility with HP technical support and management - than Business Continuity Support customers.
    - Your business receives preventive and proactive support, coupled with the fastest restoration commitment available today (restore within 4 hours of your call).

Assessing Risks (cont.)

- Training
  - System administrator and other operations staff can’t be a single point of failure
  - Complicated environments require highly technical skills to manage

Acceptable Risks (cont.)

- Service Level Agreement (SLA)
- Cost justification
  - Cost of outage
  - Lost revenue
  - Affect on customer
  - Customer satisfaction
  - Loss of customers
- Contingency planning

Designing the System

- Technology
  - Fault tolerance
  - Disaster tolerance
- Management
  - Policies and Procedures
  - Documentation
- Maintenance
  - Support
  - Training

Single Point(s) of Failure

- Price vs. Availability
- Each subsequent level of availability requires substantial investment
- Cost of outage(s) helps justify the cost incurred eliminating them

Reliability

- Know your OS and kernel settings
- Be proactive not reactive
- Regularly backup the system (restore too)
- Perform regular patch management
- Monitor and log

Redundancy

- Implement standby and alternate paths where most feasible
- Use mirroring and/or RAID for root and data disks
- Cluster

Availability

- Manage the applications
- Do proactive capacity planning
- Enforce and document change management
- Manage and maintain reliable networks
- Research and tune performance
- Perform problem management

Availability (cont.)

- Establish a realistic availability goal based on a well negotiated SLA
- Keep statistics and metrics on availability to verify you meet your goal
- View availability as a business goal as well as a technical issue
- Do a periodic reality check on availability

Policies and Procedures

- Document, document, and document some more
- Let Service Level Agreements define the goal
- Understand the business drivers and regulations
- Develop and maintain business continuity plans

Policies and Procedures (cont.)

- Focus on the environment with a Configuration Control Board (CCB)
- Record a history with an enterprise change management system
- Perform problem analysis/post mortems as standard practice
- Define escalation procedures

Disaster Recovery

- Recognize DR is different for everyone
- Know your cost of downtime
- Know what is business critical, business sensitive and non-critical
- Have multiple contingency plans
  - Hot site
  - Cold site
  - Shadow site

Disaster Recovery (cont.)

- Document your environment thoroughly
  - Take pictures
  - Keep hard and soft copies on and offsite
- Document your recovery procedures
- Recognize the procedures and plans are living documents
- Rehearse regularly
- Choose correct contingency plan

References

- www.docs.hp.com/hpux/ha (HP documentation for HA and DR)
- www.availability.com (vendor neutral site)
- http://searchhp.techtarget.com/bestWebLinks/0,289521,sid6_tax286594,00.html (searchHP.com link to high availability papers)
- www.dependability.org (IEEE CS Committee on Fault Tolerant Computing)