InterWorks/2001 Technical Presentation

Elements of Highly Available 5nines Solution Architectures

Bill Gray
Hewlett Packard Company
High Availability Advanced Technology Center
Email: bgray@hp.com

InterWorks/2001 Presentation # 121
March 2001

Introduction

There is a growing need in today's society for businesses to have continuous access to information. For large businesses using mission critical applications, even short amounts of system downtime can quickly lead to losses of customers, revenue, public confidence and stock valuation, in addition to possible financial penalties and litigation. Every business should consider the question: "What is the true cost to my business when my systems go down?"

Business Requirements for High Availability

All companies have specific availability requirements for the applications they use to operate their business. Application availability requirements can vary greatly based on how critical the application is to maintaining business operation. Some applications, such as billing or financial reporting, may be considered non-critical and pose an inconvenience if they become unavailable for a short period of time during the course of a typical business day. Other applications that perform, for example, ERP (Enterprise Resource Planning) and factory work scheduling may have a more profound effect on business operation if they are not available by potentially causing greater losses in worker productivity and business revenue due to missed product shipments. In the case of mission-critical applications, such as large centralized data warehouses that service many other company systems, downtime can potentially mean shutting down an entire business when the application is unavailable.

When categorizing the level of importance for a given application within a business (e.g. from providing non-critical services to performing operations that are an absolute necessity for the business to operate), the amount of investment required to maintain its corresponding level of availability must be considered. This investment, from a funding and staffing perspective, must be made in the areas of application development, technology infrastructure, process

implementation and support services in order to achieve the level of availability required for any given business application.

Availability Continuum

As business availability needs increase, the cost of implementing highly available systems environments also increase. This relationship between availability levels with solutions and their costs can be represented as an "Availability Continuum" Hierarchy (figure 1).

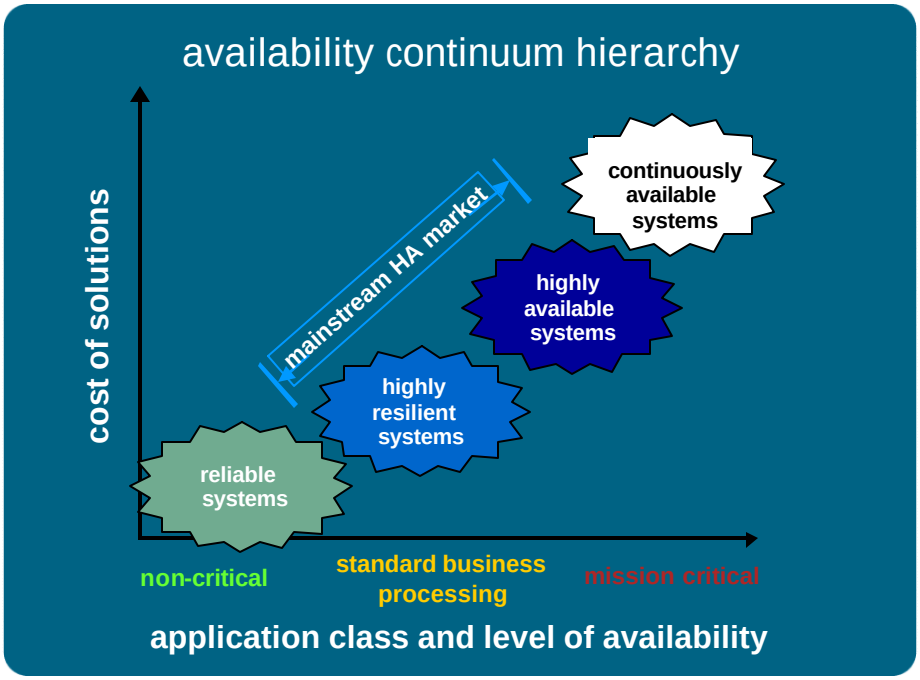

figure 1

The types of solutions available for businesses can range from:

- **Reliable systems:** These solutions utilize hardware systems, peripherals and operating systems with high mean time between failure (MTBF) rates to achieve availability levels in the range of 99% – 99.5% per year (or between 44 and 87 hours of downtime annually).
- **Highly resilient systems:** In this class of solution, availability is achieved within a single system by using internally redundant, "hot-swappable" components and data protection techniques implemented in either hardware (e.g. RAID disks) or software (e.g. OS-controlled disk mirroring), with a range of availability between 99.5% – 99.9% per year (or between 8 and 44 hours of annual downtime).

- **Highly available systems:** For this solution, externally redundant hardware is arranged in clustered system configurations and data replication techniques are used for disaster tolerance to provide from 99.9% – 99.99% availability per year (or between 52 minutes and 8 hours of downtime annually).
- **Continuously available systems**: The design for this solution carefully eliminates all single points of failure within the systems environment, allows failures to appear transparent to end-users, provides disaster tolerance with little or no data loss and focuses on IT processes to minimize planned and unplanned downtime. These types of solutions should achieve an availability level of 99.999% and above per year (equating to 5minutes or less of annual downtime).

There are a number of products from hp that provide high availability solutions for businesses along the Availability Continuum. A partial list of these products include:

Reliable Systems:

- HP 9000 L and N-class: Entry to mid-level, high performance UNIX servers designed with high MTBF rates for e-service applications. N-class systems include hot-swappable components such as internal disks and I/O cards to reduce downtime.
- Superdome: High-end UNIX servers for large-scale applications and databases. These systems are designed with many high availability features, including on-line addition and replacement of CPU, memory, PCI I/O cards, fans and power supplies, redundant power sources, error checking and correction on CPU and memory paths, and parity-protected I/O data paths.

Operating Systems:

- HP-UX: Version 11i is designed specifically for end-to-end e-services and mission-critical applications to provide scalable, 64-bit processing capability that is compatible with all HP 9000 UNIX servers. HP-UX 11i provides single-system high availability features such as support for On-Line Addition and Replacement (OLAR) of system I/O cards and components, Dynamic Processor Resilience (DPR) and Dynamic Memory Resilience (DMR) background diagnostics to detect single and double-bit errors (correcting single-bit errors and alarming on double-bit errors), dynamic tuning of kernel parameters and fast reboot times to reduce downtime.

Data Protection:

- MirrorDisk/UX: A software-based data replication product that works with the HP-UX operating system to prevent data loss due to disk failures by maintaining up to three copies of data on separate disks. This product provides high availability by allowing systems to be configured with no single point of failure for disk connections, allowing on-line backup while maintaining mirroring, and providing fast data synchronization when failures occur.
- OnlineJFS: Provides on-line management of the highly available Journaled File System (JFS), which provides a high degree of data integrity and fast file system recovery as compared to the standard UNIX file system. On-line administrative tasks include resizing, backups and defragmentation without disrupting user access to file system resources to improve availability.
- XP-family of disk arrays: A set of scalable, high-performance, highly available storage solutions for heterogeneous computing environments. The XP-family is fully compatible with all hp HA clustering solutions, and includes HA features such as RAID 0/1 and RAID 5 support for data protection, no single points of failure, the ability to perform non-disruptive upgrades, support for zero downtime backups and a battery-protected, mirrored write cache.

System Cluster Technologies and Disaster Tolerance:

- MC/ServiceGuard: A specialized, multi-computer (up to 16 nodes) clustering facility for protecting applications from a wide variety of hardware and software failures. MC/ServiceGuard monitors the health of each node within a cluster and quickly responds to failures by moving, or performing a “failover” of, operations (e.g. network connections, application services) to a functioning I/O card or node to minimize downtime.
- MetroCluster: Based on MC/ServiceGuard, MetroCluster is a metropolitan-area disaster recovery solution that minimizes downtime by providing automated failover between data centers that are up to 40km apart. Duplication of disk data between sites can be accomplished using either the hp Continuous Access XP mirrored storage solution or the EMC Symmetrix Remote Data Facility (SRDF).
- ContinentalCluster: A disaster tolerant solution that provides push-button failover between data centers without any distance limitations. Data replication between sites is performed across the Wide Area Network (WAN) to provide full restoration of operations within minutes, even in the event of a natural disaster.

System Management:

- Process Resource Manager (PRM): A facility that gives system administrators the ability to specify policies on how system resources (e.g. CPU, memory and disk) are allocated to users and applications. When used in conjunction with MC/ServiceGuard during a node failover, PRM can control how system resources are allocated on the failover system under the new load based on pre-defined resource entitlement configuration. PRM configuration changes occur dynamically during a failover and require no downtime.
- Work Load Manager (WLM): A dynamic front-end for PRM that provides system resource allocation based on administrator-defined performance goals (e.g. application response time) and workload priorities. WLM can also be used with MC/ServiceGuard to ensure that applications receive the resources they require to meet their performance goals under failover conditions.
- OpenView VantagePoint Operations (VPO): A distributed client/server product containing Intelligent Agent Clients and a Central Management Console server that provides monitoring for an entire distributed computing environment. This capability allows IT staffs to quickly identify problems before users are impacted, thus minimizing downtime.

Complete HA Solutions with hp Partners:

- Integrated Solutions: The integration of hp and its partner's technology, support and IT processes to develop complete and highly available solutions that address specific business markets. A partial list of these solutions include:
    - MC/ServiceGuard OPS Edition: Allows up to eight HP9000 servers to be configured into a scalable, highly available cluster that supports the Oracle Parallel Server (OPS) database. Its availability features include transparent recovery from LAN failures, data integrity protection and application failover.
    - Oracle Parallel Fail Safe (OPFS): A special highly available configuration of Oracle Parallel Server and MC/ServiceGuard OPS Edition that was jointly developed by hp and Oracle to provide an active/standby cluster configuration for fast database failover times.
- Reference Architectures: Repeatable, customizable, and highly reliable system architectures that have been developed to reduce time to market when implementing specific solution areas, such as ERP (Enterprise Resource Planning), CRM (Customer Relationship Management) and B2C (Business to Consumer) applications.

- 5nines Design Center: A dedicated high availability lab where hp and its partners work together with customers to design systems solutions that can achieve the highest levels of availability.

These products are all parts of hp high availability technology solutions that can meet the availability needs of businesses today (figure 2). Many of these products are used as basic building blocks for the 5nines Design Center. The entire portfolio of hp high availability products make up the “hp 5nines:5minutes Availability Continuum”. The goal of the continuum is to provide customers with choices to address their availability needs and meet their end-to-end business requirements.

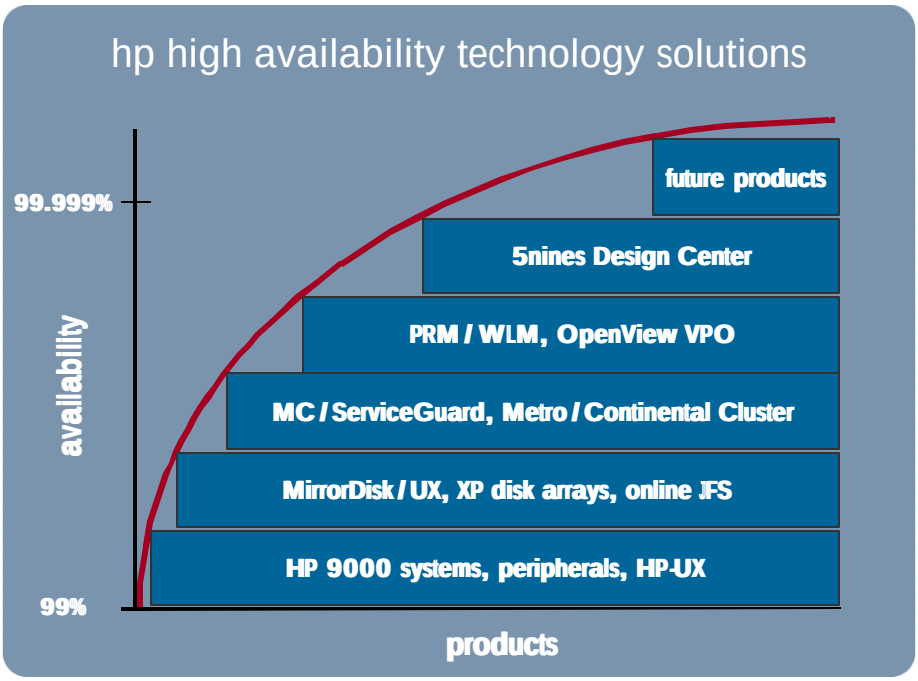

figure 2

From a cost perspective, implementing mission-critical, continuously available solutions can be from two to three times more expensive to initially create and continue to support compared to highly available system configurations. This is due to the cost of the additional redundant components required in addition to the costs associated with monitoring and managing environments that must be operational on a nearly continuous basis. It is important to consider weighing the cost of implementing a continuously available solution against the criticality of the application and cost of business losses incurred during periods of downtime in order to select the best solution for a given business need.

Hp 5nines Program

The hp 5nines Program was created to provide customers with an open systems solution that could achieve the vision of providing up to 99.999% availability (or no more than 5 minutes of downtime per year) for mission-critical applications. The definition of availability can differ depending on the point of view of the solutions provider (e.g. availability up to an operating system prompt) and the business (e.g. availability such that users can perform productive work). The 5nines solution is designed to provide availability from the system hardware and operating system, through database and middleware components up to the application level.

Downtime can be defined as either planned or unplanned. Planned downtime refers to periods of time when a system is unavailable for specific scheduled events such as backups, hardware, software and operating system upgrades. Unplanned downtime refers to the unavailability of a system due to any unscheduled events, such as hardware failures (e.g. CPU, disk, network), application failures, operator errors or power outages. The 5nines solution is designed to greatly reduce the duration of both planned and unplanned downtime events.

The 5nines Design Center is a dedicated high availability lab for the 5nines Program where hp and its partners can help businesses achieve the highest levels of availability by:

- Architecting, integrating and testing solutions based on “Best in Class” open system components from hp and its 5nines partners (e.g. HP-UX operating system, Oracle database technology, Cisco networking components)
- Providing a highly available solution stack up to the application layer
- Providing an extensible management and monitoring framework as a key element of reducing downtime within a systems environment
- Designing and documenting best practices for reduction of both planned and unplanned downtime events
- Creating a solution based not just on technology, but on IT processes for operations staffs and partner support delivery as well

Driving Needs for 5nines Solution Architectures

System downtime can come from a variety of sources. A Gartner Group study in October 1999 found that, for all the contributing factors leading to system downtime, 20% could be attributed to system hardware failures, 40% associated with application failures, and the remaining 40% associated with operator errors. Clearly, focusing on system hardware alone cannot achieve the highest level of

availability. Methods for improving applications and operations processes must also be considered.

The IDC has created the following industry definition for High Availability: Hardware systems and software systems designed to protect against component and system level failures, and when a fault or failure does occur, data is not lost and the system can recover in a reasonable amount of time. However, this definition does not take into consideration protection against the operational and process-related issues that can also lead to failures as identified in the Gartner Group study on the causes of system downtime.

Major Elements of a 5nines Solution – "Pillars of Availability"

To address all of the areas associated with system availability, hp has developed the concept of the "Three Pillars of Availability" for its high availability strategy. These pillars include: Technology Infrastructure, IT Processes and People, and Support Partnerships. Each of these pillars all work together to form a solid foundation for achieving the highest level of system availability.

Technology Infrastructure

The technology infrastructure is the backbone for any mission-critical application. Several key elements of this infrastructure include 1) the architecture of the hardware components, 2) the integration of the products used to implement the solution, and 3) how the systems environment is monitored and managed. Although each application environment will have its own unique requirements, having tested and proven pre-configured solution architectures available can provide for faster implementation and more reliable operation.

IT Processes and People

To reduce the occurrences of both planned and unplanned downtime due to operational issues, businesses must consider utilizing preventative and proactive support services. Specifically, these services should include the following capabilities:

- **Change Management:** The support service should provide an end-to-end process methodology for designing, testing, certifying and implementing any systems environment change.

- **Qualification Center:** A set of hardware, software and people resources must be provided to certify the correct operation of the solutions stack for specific hardware and software revisions.

- **Expert Board:** A staff of solutions stack hardware and software providers must be available on a regularly scheduled basis to provide guidance and consulting on all planning and design decisions.

- **Enhanced Support Services:** A staff of experts knowledgeable in the systems environment should be available 24x7 to facilitate coordination between the solution partners and to provide fast problem resolution.

Support Partnerships

Mission critical applications can be extremely complex and require elements from a number of different hardware and software providers. Each provider has it's own organizations and mechanisms for delivering support for their products, however it can be difficult for the group responsible for the overall application to manage the relationships between a variety of providers to quickly and effectively solve problems. Hp has developed close relationships with a number of industry-leading hardware and software providers (e.g. AT&T, BEA, Cisco Systems, EMC and Oracle) to create Support Partnerships that can provide a more effective level of support for systems environments.

5nines Solution Design Objectives

Although the ideas behind the "Three Pillars of Availability" strategy can be used to improve the availability of any systems environment, there are several key objectives that are specific to creating higher levels of availability in a 5nines solution architecture:

- Minimize planned and unplanned downtime
- Maximize data protection
- Maximize design flexibility

5nines Architecture Overview

Open system components provide flexibility in the design and implementation of highly available environments. Therefore, the architecture for a 5nines solution can be adapted to meet specific business environments. The following is an example of a 5nines architecture designed for a mission-critical business application that requires disaster tolerance (figure 3).

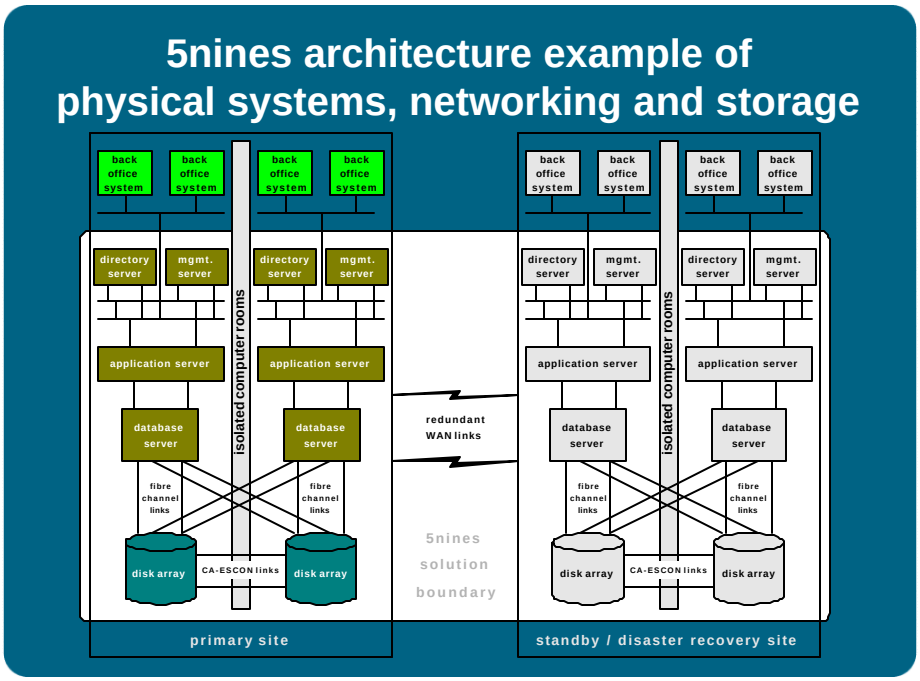

figure 3

In this example, there are two sites: a primary site and a standby / disaster recovery site. The primary site contains the application and database server(s) used to run the application. The standby site, which is geographically distant from the primary site, can assume the role of the primary site in case of a planned (e.g. system maintenance) downtime event, or during an unplanned (e.g. site-wide disaster) downtime event. The switching of operation from the primary site to the standby site is known as a “global failover”, and minimizes downtime by allowing client applications to continue execution with minimal interruption. Database records are replicated between the primary and standby sites asynchronously and in real-time over redundant and separately routed WAN links.

The systems at each site in this example are separated by a physical firewall and located in isolated computer rooms. Each room has separate power, environmental and WAN connections in case of a partial site failure (e.g. fire) that could affect one computer room while the other was still operational. Switching between systems at the same site for either planned or unplanned downtime events is known as a “local failover”. This capability, as with a global failover, also minimizes downtime by allowing client applications to continue execution with minimal interruption. Both local and global failover operations are performed using functionality provided by MC/ServiceGuard clustering.

The systems at both sites are configured as an Application Cluster and a Management Cluster. The Application Cluster supports operation of the

application and database systems, and the Management Cluster supports the management and monitoring functions of the Application Clusters at both sites.

A system within the cluster that is currently executing the application is called an “Active Node”, and a system within the cluster that handle a failover condition are called an “Inactive Node”. Active nodes within the Application and Management clusters can function independently in any computer room (e.g. the active Application Cluster node can operate in computer room #1 while the active node of the Management Cluster operates in computer room #2). MC/ServiceGuard is used on all nodes within the 5nines environment to perform the following functions:

- Startup and shutdown of the application and management clusters
- Detecting failed nodes and preventing “split brain” scenarios
- Performing automated failover operations
- Detecting and automatically switching over locally failed LAN cards

For this example, highly reliable disk arrays are located in each computer room, with one array in an active mode and the other in an inactive mode. The arrays are kept synchronized at a local level using redundant Continuous Access, Enterprise Systems Connection (CA-ESCON) links. In case of a local disk array failure, the active array in one computer room can be switched to the standby array in the other room. This capability is known as a “disk switch”. The active disk array can be in either of the computer rooms and functions independently from whichever computer room has the active application and management cluster nodes.

The systems described in the previous paragraphs fall within the 5nines solution boundary, which provides protection against hardware failures and software failures up to the application level. Systems outside the boundary can access the systems within the 5nines solution through a relocatable IP address on the active application server that can be moved to a standby application server whenever a failover operation is performed.

5nines Technology Foundation Components

The following components serve as the foundation for the hp 5nines solution:

Systems

Systems that can be used within a 5nines solution include the HP9000 L / N-Class systems and all models of the Superdome family. The HP-UX Operating System version 11.0 is required at a minimum for HP9000 L and N-Class

systems, and version 11i for Superdome systems. The characteristics of these systems include high performance, high Mean Time Between Failure (MTBF) rates, and hot-swappable components. In addition, the Superdome family of systems was designed specifically with redundant components to eliminate single points of failure. The HP-UX operating system provides both performance and scalability for supporting mission-critical applications within the 5nines environment.

Storage

Highly available, high-end disk arrays are a requirement for the 5nines solution. The disk arrays should have no single points of failure, can be upgraded without disturbing normal operation and have a variety of RAID protection levels available to optimize storage performance. Hp XP512 and EMC disk arrays are supported for hp's 5nines Solution.

Database and Data Replication

The current hp 5nines architecture uses the Oracle Parallel Server (OPS) version 8.1 as its database management system. The application clusters within the 5nines solution are configured to run OPS in a primary mode when executing at the primary site and to run OPS in a managed recovery mode at the standby site for logical data replication. This configuration provides fast database recovery when the database switches modes during a planned or unplanned global failover. The redundant disk arrays at each site perform physical data replication via CA-ESCON links to maintain data synchronization.

The management clusters in the 5nines solution also use an Oracle database as part of the OpenView VantagePoint/Operations (VPO) management and monitoring solution for the 5nines environment. VPO message forwarding performs replication of database records between the primary and standby site management clusters.

Networking

Redundant networking components (e.g. link cards, switches, routers) are used extensively throughout the 5nines solution. Multiple LAN connections in each application and management cluster node are configured to provide network traffic separation and low latency, in addition to eliminating single points of failure (SPOF) for LAN links. Networking equipment from Cisco has the capability of being controlled and monitored directly by the 5nines VPO management and monitoring solution.

Messaging Infrastructure

Client/server communication mechanisms within a 5nines environment will vary depending on the type of application used. The BEA Tuxedo transaction monitor has been integrated within the hp 5nines solution to provide:

- A separation of business logic from client applications
- An RPC communication mechanism between the client application and the database

As an optional part of an hp 5nines solution, the BEA Tuxedo transaction monitor is used to provide reliable RPC (Remote Procedure Call) communications in the case of a local or global failover by coding client applications to retry their RPCs. Tuxedo inter-domain communication is also configured within the 5nines environment to provide a relocatable IP address for the Tuxedo gateway when an application cluster active node is failed over from the primary to the standby site. Tuxedo server programs can access the Oracle database in the 5nines environment by linking libraries from the Oracle OCI library with the main Tuxedo server program that performs RPC.

It is possible to use other middleware products to provide an inter-process communication infrastructure for client/server applications that will communicate after local or global failovers. A detailed investigation of a particular middleware product would be required to determine the best way to integrate the product within the 5nines environment.

Monitoring and Management

The ability to easily monitor and manage the hardware and software components of a 5nines solution is critical for maintaining the environment and reducing both planned and unplanned downtime. Overall management of the solution environment allows maintenance operations and other scheduled activities to be performed in such a way as to minimize their impact on system users and reduce planned downtime. Extensive monitoring of the solution environment provides faster identification of component failures so that corrective actions can be performed quickly and efficiently, thus minimizing unplanned downtime.

The hp 5nines solution uses the hp OpenView VantagePoint/Operations (VPO) product as a framework for:

- **Operational Monitoring**
    - A management console GUI for monitoring key elements of the 5nines environment (e.g. systems, disk arrays, network

components, database, middleware, application clusters and VPO itself)

  - A variety of “views” available to represent the status of the environment, provide service-level reporting and to aid in root-cause analysis of failures
  - A repository for recording all events occurring within the environment

- **Management**
  - Operator commands for controlling the environment (e.g. startup, shutdown, local failover, global failover, disk switch)
  - Customization for automated and semi-automated recovery actions

The management cluster within the 5nines environment has the ability to switch roles between the primary and standby sites regardless of the location of the active application cluster. This capability provides the benefit of “follow the sun” operations by allowing geographically dispersed data centers to each manage the entire 5nines environment during normal workday hours.

VPO uses Smart Plug-Ins (SPIs) to automatically monitor the availability and performance of devices or processes within the 5nines environment. Any warnings or problems detected by the SPIs are automatically reported to the VPO management console. This provides faster analysis of failures and real-time service level tracking.

5nines Solution Stack

The combination of the concepts behind the “Pillars of Availability” and the 5nines technology foundation components form the 5nines solutions stack, which is designed to provide the highest degree of availability up to the level of the application it supports (figure 4). The 5nines solution stack includes base HA components (network elements, mass storage, HP-UX operating system, MC/ServiceGuard) and extended technology components (5nines-specific software and OPS) all being monitored and controlled by the 5nines control and management components that can utilize tools provided by solution partners. The solution stack is pre-tested and certified by a 5nines Qualification Center to ensure that the stack functions correctly for a given set of component revision levels. When the solution stack is placed into operation, 5nines support provides technical resources with specific systems environment expertise to quickly resolve problems.

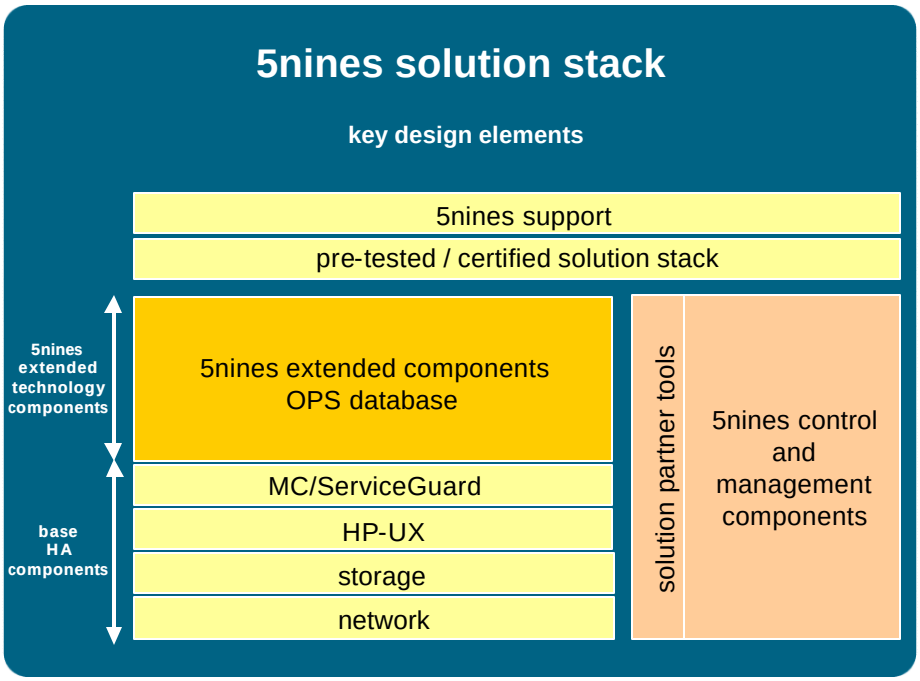

figure 4

Operational Characteristics of a 5nines Solution

The following is a list of possible failover operations that a 5nines environment can perform:

- **Planned failover:** A failover that is operator-initiated, planned and performed "gracefully" to minimize data loss and reduce downtime by allowing application clients to perform work on standby systems
    - **Local** – Failing over from an active node to an inactive node at the local site
    - **Global** – Failing over from the primary site to the remote site

- **Unplanned failover:** An immediate failover, either operator-initiated or automatically performed, to maintain service for application clients due to a detected hardware or software failure
    - **Local** – A "quick" failover to an active node within the local site
    - **Global** – A failover from the primary site to the standby site occurring due to an identified failure that effects the entire primary site

Note that all failovers can be triggered both manually (operator-initiated) and automatically (5nines-initiated) using the 5nines operational monitoring and

management framework. In the preceding 5nines solutions example, a disk array failover, or disk switch, from the active array to an inactive array is also possible in response to a disk array failure. This flexibility in failover operations allows the 5nines solution to respond to any hardware or software upgrade or failure situation at either site in order to minimize planned and unplanned downtime.

Summary of 5nines Solution Advantages

The goal of a 5nines solution is to deliver a mission-critical computing environment using a flexible hardware and software architecture with open systems components that is capable of providing maximum uptime by minimizing both planned and unplanned downtime to less than 5 minutes per year.

The following is a summary of how the 5nines solution achieves its goal:

- **Reduction in planned downtime**
    - A majority of hardware and software upgrades can be performed on a rotating basis allowing the application to execute on active nodes while upgrading inactive nodes [note there are several conditions where certain maintenance operations (e.g. application and major database upgrades) cannot be performed within the 5 minutes per year goal using today's technology]
- **Reduction in unplanned downtime**
    - The solution architecture is designed using highly available hardware, software and redundant components
    - An integrated monitoring and management framework is used for fast detection, analysis of, and reaction to hardware and software failures
    - OPS is utilized in an active/managed recovery configuration to minimize database recovery time
- **Increased Data Protection**
    - Data is replicated both at the local and standby site to preserve data integrity under all failover conditions
- **Increased Design Flexibility**
    - "Best in Class" open systems components are used to architect solutions that can meet specific business availability needs

Conclusion

There are a variety of solutions available today that can be employed to provide higher levels of availability to any business computing environment. The costs

associated with system downtime for a business must be carefully analyzed to determine the appropriate level of investment required to implement a cost-effective, high availability solution that will meet business needs. For mission-critical business environments, a 5nines solution based on the fundamental design concepts of technology infrastructure, IT processes and support partnerships can be implemented today to achieve the highest levels of availability.