

InterWorks/2001
technical
presentation #121

elements of highly
available 5nines
solutions architectures

Bill Gray
hp High Availability Advanced
Technology Center

March 2001

agenda

- business requirements for high availability
- availability continuum
- high availability solution types
- hp 5nines program and design center
- 5nines solution elements
“Pillars of Availability”
- 5nines architecture example
- 5nines technology foundation components and solution stack
- summary of 5nines solution advantages
- conclusion

business requirements for high availability

there is a growing need for businesses to have continuous access to information

for mission-critical applications, short amounts of system downtime can lead to:

- loss of
 - customers
 - revenue
 - public confidence
 - stock valuation
- financial penalties
- litigation

“What is the true cost to my business when my systems go down?”

business requirements for high availability

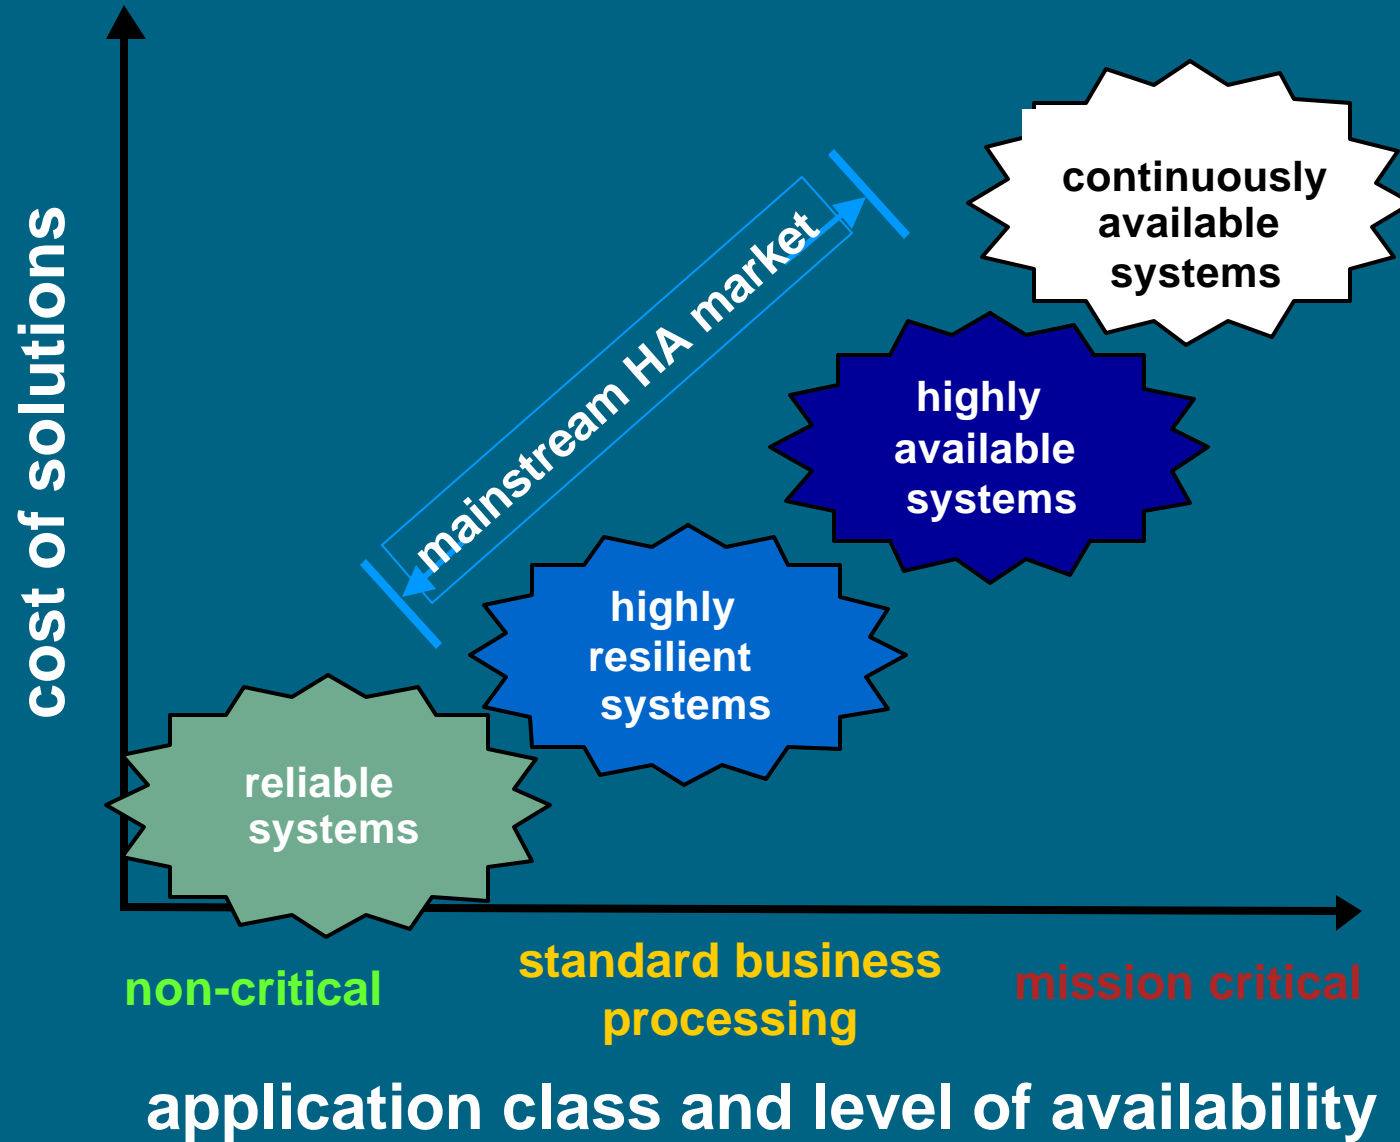
all companies have specific availability requirements for the applications used to operate their business:

- non-critical
- standard business processing
- mission-critical

investments must be made to achieve the level of application availability required in the areas of:

- application development
- technology infrastructure
- process implementation
- support services

availability continuum hierarchy



high availability solution types

- **reliable systems**

- utilize hardware systems, peripherals and operating systems with high Mean Time Between Failure (MTBF) rates
- availability levels range between 99.0 -99.5% annually
 - equates to between 44 and 87 hours of downtime per year

- **highly resilient systems**

- employs single-system HA features (internally redundant, hot-swappable components) and data protection techniques (hardware – RAID disks, software – disk mirroring)
- availability levels range between 99.5 -99.9% annually
 - equates to between 8 and 44 hours of downtime per year

high availability solution types

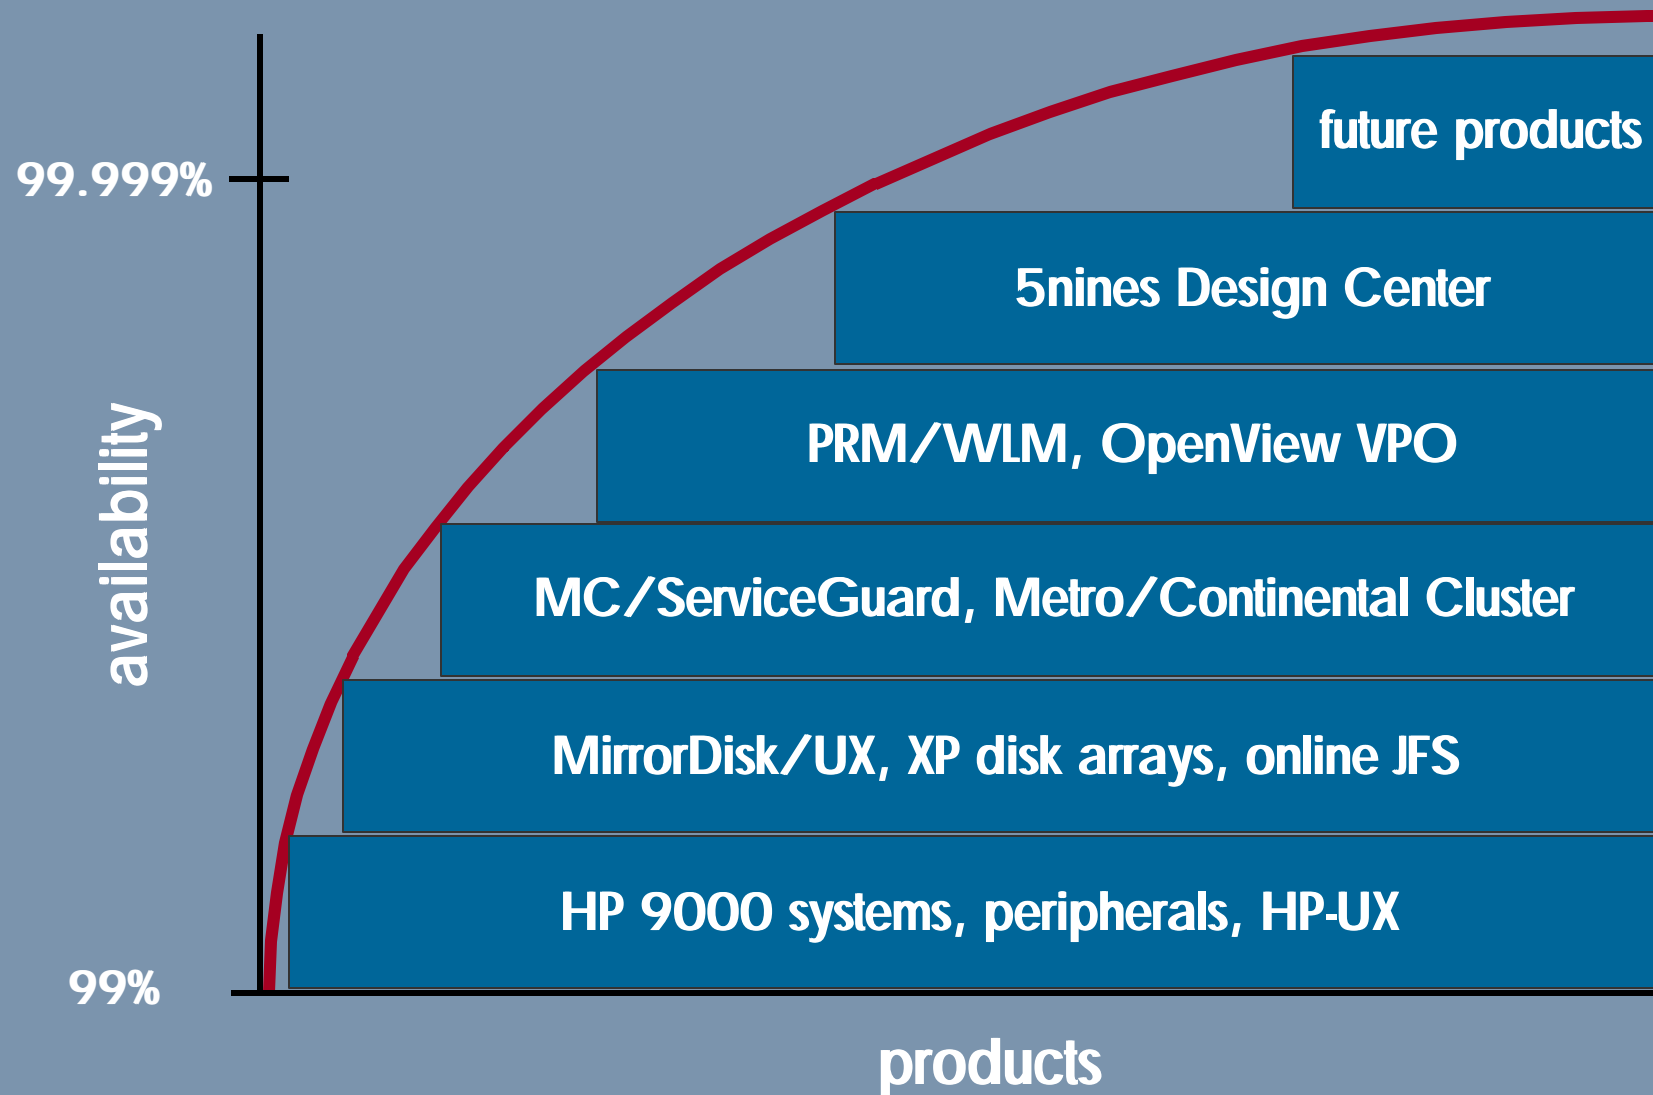
- **highly available systems**

- externally redundant hardware arranged in clustered system configurations with multi-site data replication for disaster tolerance
- availability levels range between 99.9 - 99.99% annually
 - equates to between 52 minutes and 8 hours of downtime per year

- **continuously available systems**

- solution design eliminates all single points of failure (SPOFs) within the systems environment, provides transparent failure handling for users, includes disaster tolerant features, focuses on IT processes
- availability levels range between 99.999% and above annually
 - equates to 5minutes or less of downtime per year

hp high availability technology solutions



hp 5nines program

a program to provide customers with an open systems solution that could achieve the vision of providing up to 99.999% availability for mission-critical applications

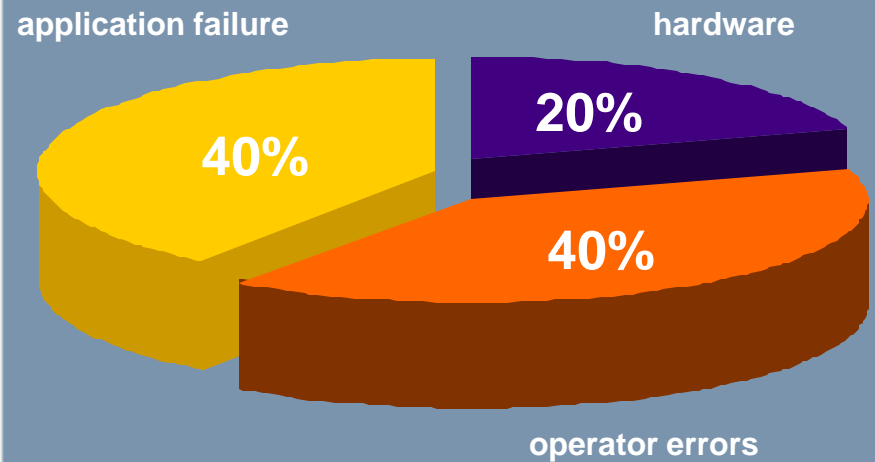
- availability is defined from the systems hardware, through the OS, database and middleware components up to the application layer
- the definition of downtime includes both planned (e.g. backups, upgrades) and unplanned (e.g. h/w or s/w failures, operator errors, power outages) events

5nines Design Center

a dedicated high availability lab for the 5nines program where hp and its partners can help customers achieve the highest levels of availability

- architecting, integrating and testing solutions using “Best in Class” open system components
- providing a highly available solutions stack up to the applications layer
- solutions are based on technology, IT processes and partner support delivery

driving needs for 5nines solution architectures




source: GartnerGroup
October 1999

high availability:

Hardware and software systems designed to protect against component and system level failures, and when a fault or failure occurs, data is not lost and the system can recover in a reasonable amount of time (IDC definition)

Question: How do you protect against operational and process-related issues that lead to system failures?

5nines solution elements “Pillars of Availability”



technology infrastructure

- architecture
- product integration
- integrated management

IT processes and people

- change management
- qualification center
- expert board
- enhanced support services

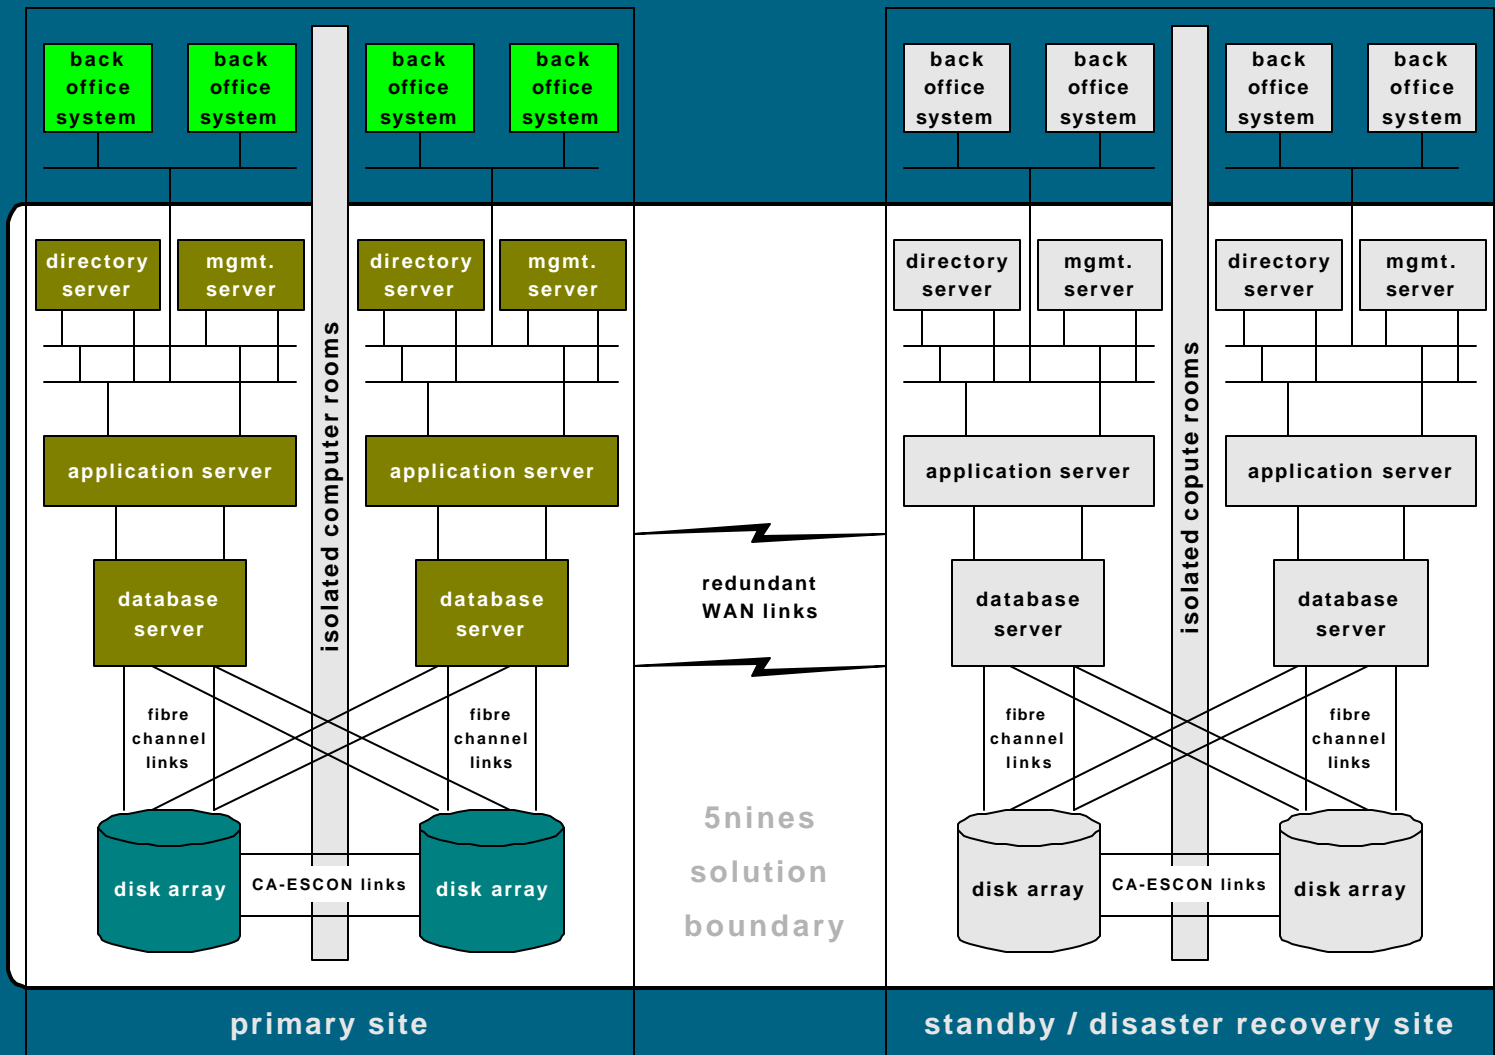
support partnerships

- partners list includes AT&T, BEA, Cisco, EMC, Oracle
- single point of contact
- integrated proactive support
- joint consulting

5nines solution design objectives

- minimize planned and unplanned downtime
- maximize data protection
- maximize design flexibility

5nines architecture example of physical systems, networking and storage



Key points of architecture example

- **data centers**

- primary and standby sites geographically separated for disaster tolerance
- two computer rooms with physical firewalls at each site
 - Each computer room has separate power, cooling and WAN connections

- **system configurations**

- systems for each site are configured as Application and Management clusters using MC/ServiceGuard, which provides
 - cluster startup and shutdown
 - detection of failed nodes and prevention of “split brain” scenarios
 - performing failover operations
 - detecting and automatically switching over failed LAN cards
- Application and Management clusters can function independently in any computer room at either site

Key points of architecture example

- **data protection**

- each site has two disk arrays in each computer room that are synchronized locally using CA-ESCON links
- database records are replicated between sites asynchronously over redundant and separately routed WAN links

- **operational modes**

- operational Application and Management services can move in response to planned and unplanned events
 - Local Failover: switching between nodes at the same site
 - Global Failover: switching between clusters at the primary and standby sites
- systems outside the 5nines solution boundary access applications protected by the 5nines solution using relocatable IP addresses

5nines technology foundation components systems

- HP9000 L / N-Class systems (minimum HP-UX version 11.0)
 - high MTBF rates
 - hot-swappable components (N-Class only)
- Superdome systems (HP-UX version 11i)
 - high MTBF rates
 - On-Line Addition and Replacement (OLAR) of CPU, memory and I/O
 - redundant components eliminate SPOFs
 - error checking and correction on CPU and memory paths
 - parity-protected I/O data paths
- HP-UX operating system
 - OLAR support
 - background diagnostics
 - dynamic kernel parameter tuning
 - fast reboot

5nines technology foundation components storage and database

- hp XP512 disk arrays
 - no SPOFs
 - RAID 0/1 and RAID 5 support
 - non-disruptive upgrades
- support for EMC disk arrays
- OPS (Oracle Parallel Server) database
 - executes in Primary mode at primary site
 - executes in Managed Recovery mode at standby site
 - provides fast database recovery during planned or unplanned global failover

5nines technology foundation components

networking and messaging infrastructure

- extensive use of redundant network components (e.g. link cards, switches, routers)
- multiple LAN connections in each node are configured to provide
 - elimination of connection SPOFs
 - traffic separation
 - low latency
- Cisco network equipment can be directly monitored and controlled by the 5nines solution
- BEA Tuxedo transaction monitor integration (optional)
- integration of other middleware IPC products into the 5nines solution is possible

5nines technology foundation components

monitoring and management

Hp OpenView
VantagePoint/Operations (VPO)
provides a framework for:

- Operational Monitoring
 - management console for monitoring elements of the 5nines environment (e.g. systems, disk arrays, network components, database, middleware, applications, VPO itself)
- Management
 - operator commands for controlling the environment (e.g. startup, shutdown, local and global failover)
 - customization for automated and semi-automated recovery actions

5nines technology foundation components

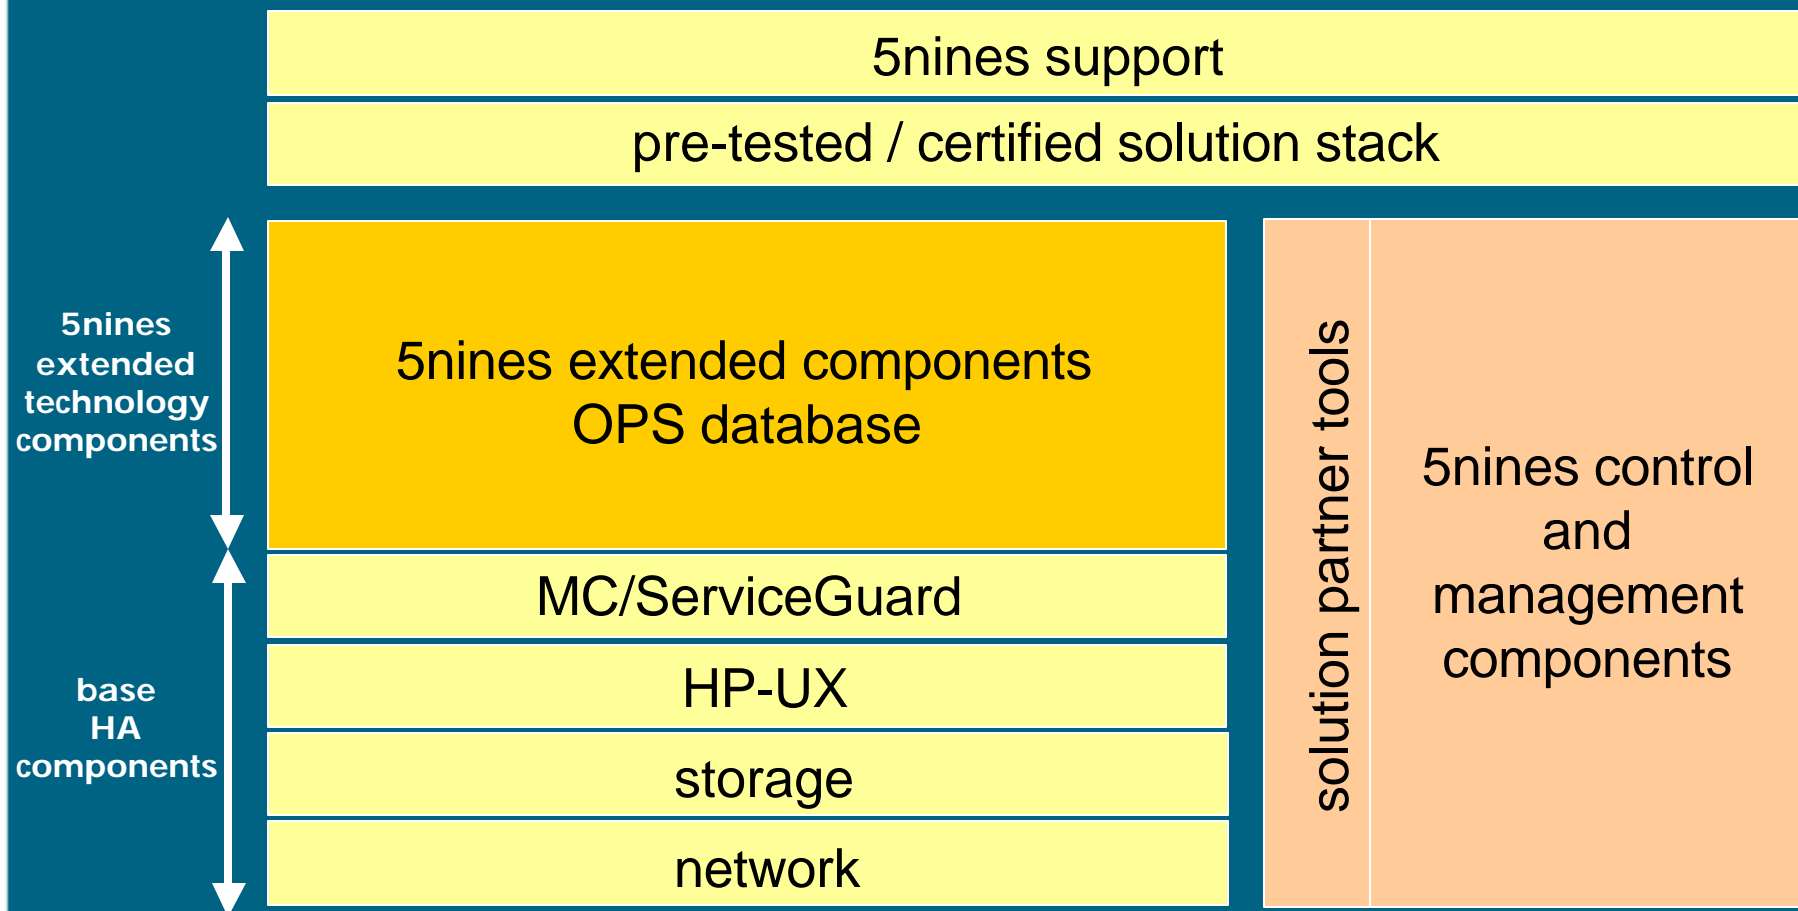
monitoring and management

additional functionality of VPO in the 5nines solution includes:

- active management cluster can switch roles between the primary and standby sites for “follow the sun” operations
- VPO message forwarding between sites
- VPO uses Smart Plug-Ins (SPIs) to automatically monitor the availability and performance of devices and processes within the 5nines environment

5nines solution stack

key design elements



summary of 5nines solution advantages

• reduction in planned downtime

- a majority of hardware and software upgrades can be performed on a rotating basis
 - applications execute on active nodes while inactive nodes are upgraded
 - note that some maintenance activities (e.g. application and major database upgrades) cannot be performed within the 5 minutes downtime per year goal using today's technology

• reduction in unplanned downtime

- environment architecture is designed using highly available hardware, software and redundant components
- integrated monitoring and management framework is used for fast detection, analysis of and reaction to hardware and software failures
- OPS is used in an active / managed recovery configuration to minimize database recovery time

summary of 5nines solution advantages

- **increased data protection**

- data replicated both at the local and standby sites to preserve data integrity under all failover conditions (planned and unplanned)
 - local site physical replication using redundant CA-ESCON links
 - remote site asynchronous replication over redundant WAN links

- **increased design flexibility**

- “Best in Class” open systems components are used to architect solutions that meet specific business availability needs

conclusion

there are a variety of solutions that can be used to provide higher levels of availability to any business computing environment

business costs associated with system downtime must be 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
- support partnerships

can be implemented today to achieve the highest levels of availability