

Making Smart Investments to Reduce Unplanned Downtime

Unplanned application downtime causes havoc and great expense. Conventional vendor wisdom focuses on redundancy to improve availability. Redundancy, however, solves just 20 percent of the problem.

Core Topic

Software Infrastructure: High Availability and Continuous Operations

Key Issue

How will organizations justify investments in technologies, people and business processes that enable high availability and continuous operations?


Strategic Planning Assumption

Through 2003, less than one third of enterprises will successfully implement change management for the heterogeneous environment (0.8 probability).

Tactical Guideline

Enterprises that lack effective production change control processes will not achieve 98 percent or greater application availability levels.

Figure 1
Causes of Unplanned Downtime



Source: GartnerGroup

Note 1

Incident Management

Evaluating downtime requires an understanding of the incidents and the length of outage caused by each. Technology failures and disasters occur less often than other types of failures (e.g., the average hardware in a PC server has a 12 to 18 month mean time between failures, and major data center disasters occur an average of once every 20 to 100 years); however, the length of the outage is often measured in hours or days, and thus, such failures can be catastrophic.

Based on extensive feedback from clients, we estimate that, on average, unplanned application downtime is caused (see Figure 1): 20 percent of the time by hardware (e.g., server and network), OSs, environmental factors (e.g., heating, cooling and power failures) and disasters; 40 percent of the time by application failures including “bugs,” performance issues or changes to applications that cause problems (including the application code itself or layered software on which the application is dependent); and 40 percent of the time by operator errors, including not performing a required operations task or performing a task incorrectly (e.g., changes made to infrastructure components that result in problems and incur unexpected downtime).

Thus, approximately 80 percent of unplanned downtime is caused by people and process issues, while the remainder is caused by technology failures and disasters. Improving availability requires a different strategy and set of investment choices for each of the three unplanned downtime categories.

Technology Failures and Disasters: Despite being just 20 percent of unplanned downtime, these types of failures can be very catastrophic and result in a significant amount of downtime per incident (see Note 1). To mitigate this risk, enterprises should take the following steps. *Monitor* components for availability (since failure identification is the first step toward resolution). This is typically done with agents or sensors. Ideally, monitoring is predictive and warns the operator or vendor of potential failures prior to their occurrence. *Buy* vendor service contracts to reduce time to repair. Many vendors offer time-to-repair commitments for increased fees. *Implement* redundancy to ensure alternate processing capabilities in the event of a catastrophic failure. Data mirroring, clustering and diesel generators are examples of redundancies that limit downtime when failures occur. In comparing potential solutions, pay particular consideration to

GartnerGroup

how transparent recovery is to end users and applications.

Note 2

Packaged Applications and IT Processes

For purchased application components, enterprises concerned with application availability should evaluate the vendor's IT processes as part of the package selection criteria and work closely with their selected vendors to improve and re-engineer the vendors' IT processes, especially those related to software and integration testing (across vendors' products) and application architecture and design.

Note 3

Change Management Process

Change management is the discipline that preserves the integrity of the production environment and reduces the risk of service degradations and interruptions caused by change. Rather than each individual or department orchestrating independent changes on desktops, networks, systems, software and facilities, a coordinated and structured approach significantly reduces unplanned downtime caused by problems that often occur when changes are made. In addition, planned downtime is reduced considerably by attaining maximum leverage for each downtime period (while considering risk). Steps in the process are the following:

1. Managing and tracking change requests
2. Approving changes while considering business and technical risk
3. Scheduling changes and change notification
4. Project management, planning, building and testing changes
5. Implementing and deploying changes
6. Assessing results (which then become input to the change approval process)

Acronym Key

- OS** Operating system
SLA Service-level agreement

Application Failures: To reduce downtime caused by application failures, enterprises should invest in improving and re-engineering IT processes (see Note 2), including the following: 1) *change management* — reduces unplanned downtime caused by inadequate planning and testing of application changes, enables a more proactive approach toward problem prevention (see Note 3); 2) *problem management* — improves problem identification, isolation and resolution, thereby reducing time to repair; 3) *configuration management* — tracks the relationships between dependent application and infrastructure components, enables better understanding of change impact and quicker fault diagnosis; 4) *application architecture and design* — reduces single points of failure, aides in problem isolation and makes application failures more transparent to users; and 5) *performance management and capacity planning* — proactively identifies current and future resource shortages impacting SLAs.

Re-engineering processes requires investing in people responsible for the process and tools to facilitate the process. By comparison, reducing technology failures through redundancies is a much simpler problem to solve than process re-engineering, since redundancy-related failures do not require significant changes in human behavior. Through 2003, less than one third of enterprises will successfully implement change management for the heterogeneous environment (0.8 probability).

Operator Error: Reducing downtime caused by operator error requires: 1) maturing IT operations to a more-process-oriented and documented approach that does not require or mandate specific, knowledgeable people be available to perform tasks, 2) hiring competent people and training operators (and vendors) on IT process and procedures, 3) automating the process wherever possible to reduce the chance for errors (e.g., using job scheduling and event management tools), and 4) improving the change and problem management processes related to IT infrastructure and facilities.

Bottom Line: Enterprises should not let infrastructure redundancy provide a false sense of availability assurance. To address the 80 percent of unplanned downtime caused by people and process failures (vs. technology failures or disasters), enterprises should invest in improving IT processes, such as change, configuration and problem management; performance and capacity planning; application architecture and design; and operator hiring and training. Investments should also be made in automation. Other downtime causes should be addressed by eliminating single points of failure through redundancy, with vendor service contracts and component monitoring.

