Availability In Products and Services

Terje Iversen
Twenty Four Seven LLC

1 Introduction

Continuous access to critical data and key information services is central to the survival of firms competing in today's e-business economy. To be marketed successfully to such firms, products must meet lofty customer expectations: they must afford the most appealing and apposite features, they must be constantly ready to function, and they must be backed by excellent customer support practices.

Successful product development companies are customer-centric and have programs, processes and tools in place to render products without defect and peerless in filling their jobs. They know that product quality, continuity of function, and customer service are the ultimate differentiators among suppliers, and that only superb products and services lead to high levels of customer satisfaction, repeat sales, and good references to help increase market share.

These points are well-known also to buyers in businesses operating in real time who understand the costs of outage of crucial systems. This awareness is reinforced by evidence like that presented in Figure 1.

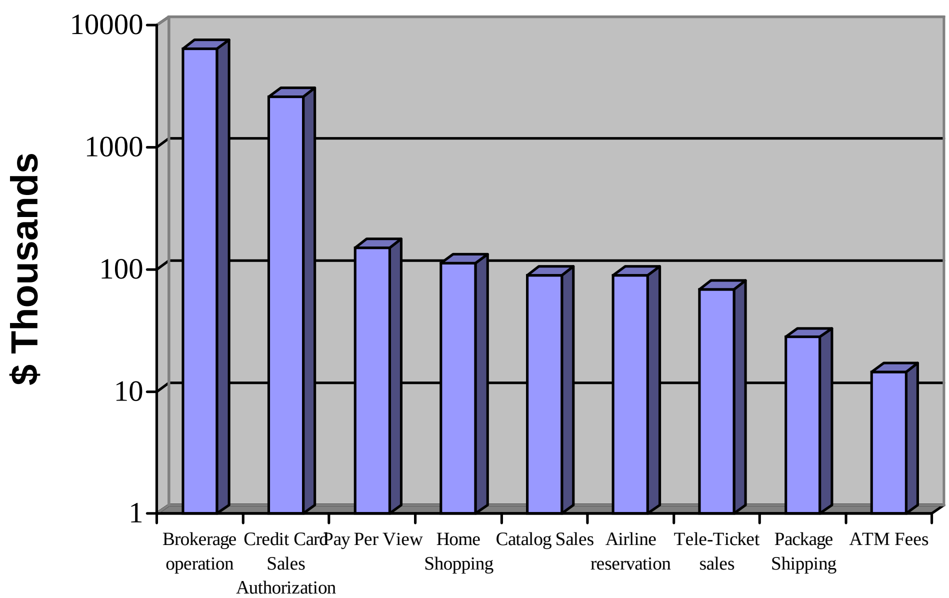

Hourly Downtime Costs by Industry

(Reference: Fibre Channel Industry Association)

Figure 1

A brokerage, the figure tells us, may lose ***$6,000,000*** for every hour its operations are disrupted because the function of one or more of the many products underpinning the operations is unavailable.

With such enormous consequences for failure, it is not surprising that customers demand ever higher likelihood that a product or service will perform as specified at any chosen time when negotiating to buy new products and services. In fact it is common to ask for, “five nines” or 99.999% probability of freedom from failure in service level agreements today.

2 High Availability

The probability that a product or service will be outage-free at any moment is termed its *availability*. High availability is achieved through a combination of reliability, repairability, and redundancy.

- *Reliability* of a product or system is the demonstrably typical duration of its uninterrupted usage and is expressed by its "Mean Time Between Failures" or MTBF.
- *Repairability* is a measure of how fast a system or product can be put back into service after a failure and is expressed by its "Mean Time To Restore" or MTTR. Implicitly, the speed of returning a failed product to operation is influenced by diagnosability of failures and by the responsiveness of repairmen.

$$\text{Availability} = \frac{MTBF}{MTBF + MTTR}$$

- *Redundancy* helps improve availability by establishing a backup or standby for selected individual components or the complete product. The effectiveness of redundancy depends on how quickly the backup or standby can be introduced into service. As an example, redundancy with instant and automatic substitution of a standby, called *fail-over,* ensures no impact on operations when the component for which redundancy is provided fails.

Table 2 shows the relationship between availability values, the average total downtime per year they imply, and the MTBF required to achieve each availability if MTTR were 2 hours.

Availability	**Avg. Downtime/yr**	**Required MTBF: MTTR = 2 hr**
99%	4 days	200 hrs or 8 da
99.9%	9 hrs	2,000 hrs or 3 mos
99.99%	1 hr	20,000 hrs or 2 yrs
99.999%	*5 min*	*200,000 hrs or 23 yrs*

Table 1

As can be seen from Table 2, to reach five nines availability the product can not be down more than five minutes a year on average; if the product's MTTR is two hours, its MTBF must be about 200,000 hours or 23 years to achieve the five-minute annual average.

2.1 Redundancy Benefits

Redundancy with automatic fail-over has a very positive effect on availability. If a component with reliability *MTBF* and repairability *MTTR* fails but has a redundant component to which the operation can be switched without interrupting the application that is running, the redundant mean time between failures, $MTBF_R$, can be expressed as follows:

$$MTBF_R = \frac{MTBF^2}{2 \bullet MTTR} \quad \text{when } MTBF >> MTTR$$

As an example a component with an *MTBF* of 1000 hours and an *MTTR* of 2 hours will have with redundancy an $MTBF_R$ of 250,000 hours. Of course, the component can itself be a subsystem or a system[1].

Availability under redundancy in a component is the fraction of time when at least one component is not in need of repair. The expected duration of simultaneous outage of two components is the average overlap of the two outages, randomly phased, and is easily seen to be one-half the *MTTR*; this we may call $MTTR_R = MTTR/2$. Availability is found by adapting the earlier equation:

$$\text{Availability} = \frac{MTBF_R}{MTBF_R + MTTR_R}$$

And in the example above,

$$\text{Availability} = \frac{250{,}000}{250{,}000 + 1} = 0.999996$$

2.2 Composite Products

From an availability standpoint, a product that includes both hardware and software can be viewed as a product with the two components in series, each with its own availability, since failure of either of the two components will make the product unavailable. A product with components in series has an overall availability that is equal to the product of the availability of the individual components.

$$\text{Availability}_{\text{total}} = \text{Availability}_{\text{hdw}} \text{ x } \text{Availability}_{\text{sw}}$$

Thus, a product with both hardware and software must have availability greater than five nines in each component to get an overall product availability of five nines. As an example, if both hardware and software have the same availability, each must have an

[1] This calculation relies on the assumption that a failed component can be repaired while its alternate is serving the production need. It yields the mean interval between incidents of simultaneous outage of the two available components.

availability of 99.9995% to have a product availability of five nines, an average down time of less than two and a half minutes a year for each component.

3 Design, Delivery And Support Of High Availability Products

Designing, delivering and supporting high availability products are enterprise-wide undertakings and are achieved through a combination of the following practices.

- Formal product development process
- Design for highest quality
- Design for manufacturability
- Design for supportability
- Project management and cross-functional project teams
- Formal support processes

Careful management of a product through its lifecycle of product planning, design and development, introduction, operation and decommission is required to achieve high quality and high availability. In the planning and design phases, quality and availability requirements are established and must forcefully govern the design of the product. During the introduction and operation phases, feedback from customers and from manufacturing and support staffs must be used to improve the product and update the processes used to design, deliver and support it.

Product failures and unavailability happen as a result of product defects, human errors and process errors. New product development or programs to improve availability in existing products must seek to eliminate the defects. But since, as we have seen, the time to restore a product after a failure so profoundly influences availability, requirements for serviceability and maintainability must be included as primary in the product design.

Product defects that lead to failures, hence compromise availability, result from errors in design and component selection, from errors built into the product in manufacturing, or from errors in deployment of the product.

Missing features and functions may lead to customer dissatisfaction and reduce the likelihood that customers will be repeat buyers and provide good references, but they don't generally lead to increase of failures. In at least two ways, though, feature lack may gravely affect *availability*: (1) if the diagnosability of problem causes is made weak or (2) if adequate tools for easy, accurate, and suitable deployment of the product or for defending the product against misuse are omitted.

According to the Gartner Group about 40% of all down time is due to human errors and process errors. Thus a careful analysis of expertise, staffing, documentation and processes must be made to determine and close the gaps that may affect availability.

3.1 Formal Product Development Process

Product development is a phased process that begins with a concept and continues with design, implementation, verification, validation and release. Successful product development organizations have a formal methodology for this process. This is not to say that companies without a formal process are unable to develop quality products. Rather, the success of an undisciplined development program is generally through the heroic efforts of a dedicated team and repeating the result depends entirely on having the same individuals available for the next project. Success that rests solely on the availability of specific persons provides no basis for continued quality and availability improvements.

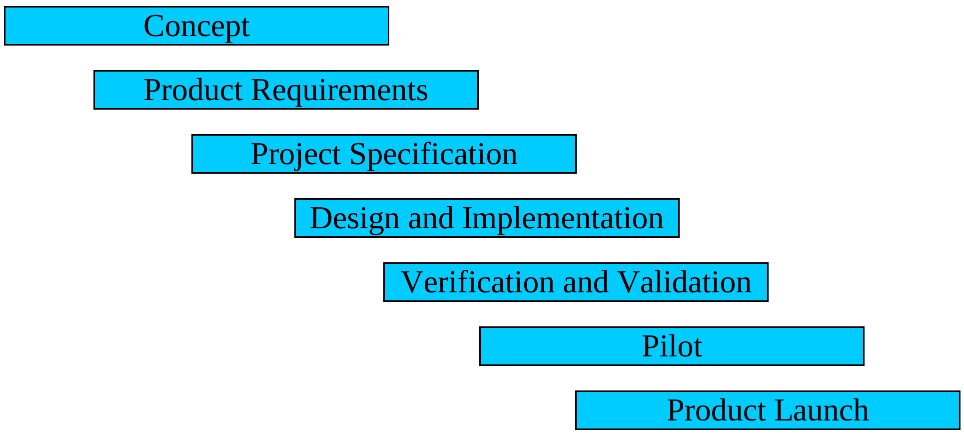

Phased Development Model

Figure 2

Figure 2 shows a typical phased development process. Successful product development organizations use similar processes, customized to fit their needs. The process must be documented including definitions of checkpoints, deliverables and exit criteria from each phase and it must be managed to make sure that it is followed and that corrections are made to remove obstacles as they occur.

Flexibility must be built into the process to accommodate minor changes as the development progresses to take advantage of knowledge gained during the development and to make sure that market changes will not make the product obsolete. To accommodate these changes, it is important that the product architecture be structured in such a way that minor changes in functionality as the project progresses can be made without requiring major changes to other parts of the system. In addition it is important that parts of the design that are uncertain or likely to change are buffered from those parts that are expected to remain stable, and that the product is partitioned so that critical elements of system performance can be demonstrated even when some of the individual modules are only partially complete.

3.2 Design for Quality

A quality product is a reliable product that meets customers' needs, satisfies their expectations on features and functions, and has the availability needed to get their work done. To ensure that these requirements are set and met, all the relevant stakeholders must be included in the design process. Primary responsibility for reliability rests with the product developers; they must, however, be supported by the whole enterprise, and by customers and suppliers as well. Customers must be involved to make sure all applicable requirements are captured in the requirements phase and suppliers must be included to ensure that quality parts to build the product will be available.

Quality assurance staff must work closely with product developers during the design process to ensure that quality and testability requirements are understood and reflected in the design and that best practices are used to remove defects from the product long before it reaches the customer. These practices include peer reviews, inspections, module testing, simulations using design tools and root cause failure analysis. The importance of eliminating defects as early as possible in the design cycle is illustrated in Figure 3. In addition to being expensive, removal of defects after the product is released results in lowered customer satisfaction.

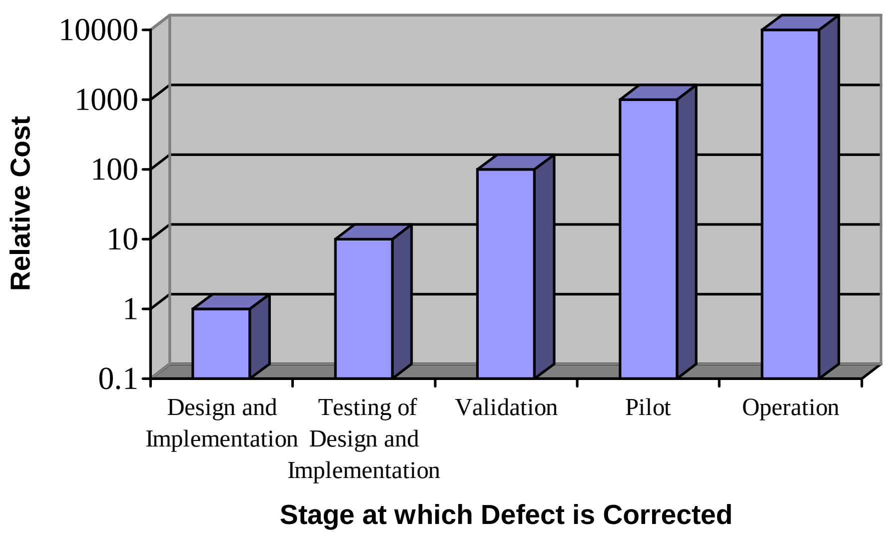

Figure 3

Root cause failure analysis is a crucial necessity when designing reliable products. It is essential to get to the root cause of every problem so that it gets resolved permanently. Often symptoms are addressed instead of causes so problems linger and improvements are not completely effective, often with negative results for the customers.

A product failure database or “knowledge base” that is accessible to the enterprise must be included for root cause analysis to be effective. This database is vital to product developers to avoid repeating the defects of past products, and to such other groups as field service, customer service, procurement and inventory control, accounting, suppliers, and manufacturing to help improve product reliability throughout the product life cycle.

3.3 Design for Manufacturability

Design for Manufacturing (DFM), a professional field in its own right, means designing to optimize the manufacturing function to assure the best quality, reliability, and customer satisfaction, and as well, to maximize efficiency of manufacture. An overriding objective is to obviate premature failures that lead to poor availability.

A product must be designed for ease of fabrication, process and assembly in order to minimize problems in the manufacturing process and avoid mistakes that build defects into the product.

All parts selected in the design process must be from reliable sources with defined and guaranteed failure rates to avoid premature failures and poor availability. Additional benefits of DFM are a short time to market and a smooth transition to manufacturing from product development.

3.4 Design for Supportability

Supportability includes not only serviceability and maintainability, but suitability for all the events—including use, upgrades, and reconfigurations—during customer ownership from installation to decommission. Design for Supportability (DFS) , another specialization, offers great benefit to availability, customer satisfaction, and the cost of delivering service, but it is often considered too late in the product development process to have those effects. In fact, surveys have shown that, on average, product support is not considered until two thirds of the development cycle is completed; by this time most of the design decisions have already been made.

Examples of supportability features that should be considered in the requirements and design phases of the product development process are:

- Redundancy and fast automatic fail-over
- Hot removal and replacement
- Automatic failure detection and reporting
- Automatic data collection from failures
- Remote maintenance
- Recovery/restart/reboot time loss
- Modular design to reduce troubleshooting and repair costs

As discussed earlier, redundancy improves reliability and if it is implemented with automatic fail-over and hot removal and replacement, it improves product availability and lessens the pressure on the field support organization to urgently remove and replace the failed component or system.

In order for automatic fail-over to occur, automatic failure detection must be implemented. If detection is also coupled to a remote maintenance facility, the field support organization will be alerted to the failure and replacement of the failed component will be scheduled with no action by the customer.

Automatic capture of failure data combined with remote access to this data makes it possible for the support organization to retrieve it for failure analysis without disturbance to the customer's operation and to incorporate the data in the enterprise knowledge base.

The time it takes to restart a system after a failure, a preventive maintenance episode, or an upgrade is often not considered in the design process and its impact is not felt before the system is in operation. This usually unhappy surprise can be avoided by considering preventive maintenance and upgrades in the design process and including requirements on restart and reboot times. If requirements to allow inline maintenance and upgrades are also included, any problem for the customer of rebooting and restarting the product can be substantially reduced or eliminated.

Principal benefits of a modular design are in cutting the time it takes to troubleshoot a problem and reducing repair costs. In a modular system, a failure can be localized to a particular module and repair only affects that module. Further, changes made on a module do not propagate trouble to other parts of the system and risk additional failures.

3.5 Project Management and Cross-Functional Project Teams

The phased product development process shown in Figure 2 requires the use of project management and project management tools to plan, monitor and control the process. A plan is developed at the beginning of the project with definitions of all the key tasks, their dependencies, and their needed completion dates.

The project is monitored to make sure progress is consistent with the plan, that the product development process is followed, that the project meets time and cost goals and produces all the deliverables specified for each phase. Close scrutiny shows where corrections are needed to break down barriers and remove obstacles as they occur.

Risk identification, quantification and control should be active. Continued focus on risk helps to identify potential problems before they occur and can ensure that risk handling activities are planned and used when needed. Risk management lessens unfavorable impacts on cost and schedule, and avoids compromise of overall objectives.

Many enterprises today consider cross-functional product development teams to be an integral strategy for success. The cross-functional team or integrated product team should be composed of representatives of all groups that have stakes in the success of the product development effort. The strategy leads to increased job satisfaction and better cooperation among departments. It is a certain way to ensure that the overall quality and reliability goals set at the start of the project are met, and that all the requirements for testability, manufacturability, serviceability and maintainability are realized in the design of the product.

Additional benefits of project management and cross-functional product development teams, are a shortened development cycle, a shortened time to market, and lowered development cost.

3.6 Formal Support Processes

Formal documentation of support processes is necessary for process management, process repeatability, and process improvement. Important support processes are configuration management, process and product quality assurance or QA, and customer support.

Configuration management establishes and maintains the integrity of work products by:

- *Identifying the configuration of selected work products that make up the baselines at given points in time*
- *Controlling changes to configuration items*
- *Building or providing specifications to build work products according to configurations under management*
- *Maintaining the integrity of baselines*
- *Providing accurate status and current configuration data to developers, end users, and customers*

Work products placed under configuration management include the products that are delivered to the customer, designated internal work products, acquired products, tools, and other items that are used in creating and describing the work products. [1]

Process and product QA provides the enterprise with objective insight into processes and associated work products by:

- O*bjective evaluation of performed processes, work products, and services against process descriptions, standards, and procedures*
- *Identifying and documenting non-compliance issues*
- *Providing feedback to the enterprise on the results of quality assurance activities*
- *Ensuring that non-compliance is addressed* [1]

The objective of the QA processes is to support the delivery of high-quality products and services by providing the enterprise at all levels with visibility and feedback into processes and associated work products throughout the product life cycle. To effectively carry out this responsibility the QA organization must be independent of the project organization in order to establish objectivity. QA should be involved from the beginning of a product development project to help establish plans, standards, processes, and procedures to satisfy the requirements of the product development project and organizational policies.

Customer support processes cover the phases of the product lifecycle from installation, through use and decommissioning. The objectives of these processes go far beyond resolving problems as fast and as cheaply as possible. The processes include activities such as:

- Customer relationship services
- Technical support services
- Problem escalation

- Problem resolution
- Logistics
- Training
- Documentation

The overarching goal of support is to maximize the benefit to the customer during the period of ownership to achieve high levels of customer satisfaction in order to ensure customer retention, repeat buying and positive references to help raise market share.

4 Conclusions

Enterprises that are successful in designing, delivering and supporting high availability products and services use a combination of processes and best practices to manage the product life cycle. They have a strong reliability culture that motivates everyone in the enterprise. The processes and best practices are formal, documented, and institutionalized. Reliability goals are set and made visible, and progress toward these goals is tracked and published. Processes and best practices are continuously audited and updated to ensure that customer satisfaction and product quality goals are met and that customers and suppliers are included in all programs to improve product availability.

5 References

1. Capability Maturity Model Integration (CMMI^{SM}) Version 1.1, Continuous Representation (CMU/SEI-2002-TR-011, ESC-TR-2002-011) Carnegie Mellon Software Engineering Institute
2. Pinto, Jeffery K. “Project Management 2002” *IEEE Engineering Management Review* 31(1) (First Quarter 2003): pp. 92-106
3. Goffin, K. “Evaluating Customer Support During New Product Development-An Exploratory Study.” *Journal of Product Innovation Management*, 15(1), (January 1998): pp. 42-56
4. Van Bennekom, F. and Goffin, K. : Using Customer Support’s Knowledge Base to Enhance New Product Development.” *Service Operations Management Association Conference Proceedings*, 1999
5. Cisco Systems White Paper, “IP Telephony: The Five Nines Story.”