

A Foundation for System Availability

How Nimble Storage leverages the power of big data and cloud technology to achieve greater than 99.999 percent availability

As technology advances, the effects of unplanned downtime increase along with the expectations of high availability and resiliency, particularly in larger enterprise environments. Recent advancements in storage lifecycle management, including cloud-based technology and deep data analytics, are contributing mightily to achieving these goals. This paper discusses how Nimble Storage has been on the forefront of these technology evolutions, and how the current Nimble installed base of storage systems has achieved the important benchmark of 99.999 percent (or “five nines”) annual uptime in just over three years of shipping its systems.

Introduction

The importance of reliability and availability of storage systems in today's larger environments continues to grow over time. These requirements now apply even to non-mission-critical but important applications as well. One of the standard metrics for measuring availability is system uptime, and the ultimate benchmark for success is achieving 99.999 percent annual uptime, or "five nines."

Several factors go into achieving the level of availability required by today's demanding environments. Nimble Storage, a leader in flash-optimized hybrid storage, is built on the patented Cache Accelerated Sequential Layout architecture (CASL™). CASL is designed from the ground up to leverage both the power and performance of flash technology as well as the cost-effective capacity of hard disk drives. The flexibility in designing a new storage platform also allows CASL to address availability concerns as well, offering capabilities such as an active/standby controller architecture that allows for completely non-disruptive upgrades and virtually zero performance impact due to controller failure, robust protection against data loss with fast multi-parity RAID, and overall data reliability and integrity. Fault tolerance has also been built into the hardware platform across the portfolio, leveraging redundant components with no single point of failure. In addition to the design elements that contribute to overall availability, Nimble has taken additional steps that capitalize on recent developments in deep data analytics and cloud-based management that contribute to the overall uptime across Nimble's customer base. We will focus on these latter developments in this paper.

First, let's examine Nimble's approach to data analytics as it relates to both storage management and support. Nimble has developed an entirely unique and innovative approach to storage lifecycle management that leverages the power of data sciences. The result is Nimble Storage InfoSight, which integrates, automates, and substantially simplifies administrative tasks, ensuring the optimal health of all Nimble Storage arrays. Built on powerful deep data analytics technologies, the centralized InfoSight Engine monitors all Nimble Storage assets collectively from the cloud. The engine analyzes millions of data points collected from each array on a daily basis across the Nimble installed base to build complete insight into overall storage system health. Nimble also employs an aggressive and proactive support model, where all unplanned downtime events are investigated to find root causes. Fixes are then pushed out, not only to the affected customer, but also to customers with systems at risk of the experiencing the same issue. The combined effect of these factors is that overall availability increases over time and the likelihood of a system experiencing a downtime event becomes increasingly rare.

This unique approach has led to a combined observed uptime of greater than five nines across the Nimble customer base, a significant achievement in slightly more than three years since delivering the first Nimble storage system.

The following sections explore the methodology of how Nimble has measured and observed total system uptime, and also offer several other statistical views into the resiliency of Nimble Storage arrays in general.

Methodology

Nimble Storage arrays are equipped with sophisticated logging processes that provide a rich set of historical data at frequent intervals. This data provides a complete history, with the duration of downtime events known to the second, from which system availability can be calculated and any downtime events can be classified. Customers can choose whether or not to have this data sent back to Nimble. We performed analysis on thousands of arrays that were reporting data to Nimble during the time period from July 2012 through November 2013. This data set totals 1,750 years of combined system time.

Internal test systems were excluded from analysis. It is important to note that all of the results presented in this paper are not based on artificial data or design principles, but rather represent observed values from Nimble systems in production.

First, we divide an array's history time into several meaningful categories. A system may be up, i.e., available for normal use, or down. Downtime is further categorized into unplanned, planned, or environmental events. Planned downtime consists of events where the customer shuts down the array for a datacenter move or another procedure that results in a period where the array is unavailable. (Note: Upgrades to Nimble arrays are non-disruptive.) Environmental downtime refers to situations where the array is unavailable for reasons that have nothing to do with the array itself, such as power outages or network problems. Unplanned downtime is downtime that results from a problem with the array itself and is the focus of our analysis. We use a standard definition of availability, the ratio of uptime to the sum of uptime and unplanned downtime.

Analysis

Empirical results

The breakdown of time by category in our data is:

| | |
|------------------------|-----------|
| Uptime | 99.41258% |
| Planned downtime | 0.57925% |
| Environmental downtime | 0.00789% |
| Unplanned downtime | 0.00029% |

This gives an observed availability value of 99.9997% across the Nimble installed base of thousands of systems.

Estimating Confidence

While our observed data shows greater than 'five nines' availability, even if we only consider the minority of systems that experienced some unplanned downtime, it is prudent to provide a measure of confidence in our results. Some downtime events are caused by randomly occurring issues and all data is influenced by chance. There are also customers that have elected to not provide us with data. A statistical analysis is required to account for these sources of uncertainty and incorporate them our estimate of overall availability. The presence of complete historical data removes some complications that are typical in this type of analysis, such as the inaccuracies in the length of downtime events or the possibility of undetectable events.

The standard approach is to construct models that describe the probability of an array experiencing downtime events and the time associated with each event. With these models one can investigate the relationship between chance and availability and determine the probability that overall availability is actually below a specified level, in our case 99.999 percent. This is a well-studied field in statistics and many techniques are available. However, the scarcity of downtime events and the lack of variance in downtime lengths across the Nimble installed base makes regular approaches unusable. Most downtime lengths are very short, and the few that are not are exceptional, due to the impact of Nimble Storage's approaches to design and support. Innovations in system and software design result in short downtimes when not actually eliminated, and, when significant downtimes happen, an aggressive support model often supplies fixes to at-risk customers in advance of their experiencing the issue. We are left with a number of small events that do not vary and a few large events that do not generalize.

To boost confidence, we looked at new ways of approaching the problem. Since a fundamental issue with our analysis is the lack of arrays that experience downtime we looked at the effect of exaggerating the probability that an array would have one or more unplanned downtime events. As this factor is increased and problematic systems

are overrepresented, availability slowly drops, variance increases, and we can track the increasing large confidence intervals. Eventually, we can produce a confidence interval where we can answer our original question of confidence in the observed availability. The results of this are shown in Figure 1.




Figure 1: Confidence limits for availability as the presence of downtime-experiencing systems is artificially increased

For there to be a 1 percent chance that overall availability is below 99.999 percent, problematic systems would need to be produced at a rate 40 times greater than observed. In other words, we are 99 percent confident that availability will achieve 99.999 percent or greater, even by over-representing problematic systems by a factor of 40. Though this is not a traditional confidence statement, it does give an indication of how far our observed results are beyond the 99.999 percent threshold.

Availability Over Time

Nimble’s support process and focus on availability helps increase availability over time, as data is gathered and downtime-producing issues are identified and fixed. Often, potential problems are resolved in advance of triggering a downtime event.

As seen in Figure 2, the longer a system runs without experiencing unplanned downtime, the less likely it is to experience it in the future.

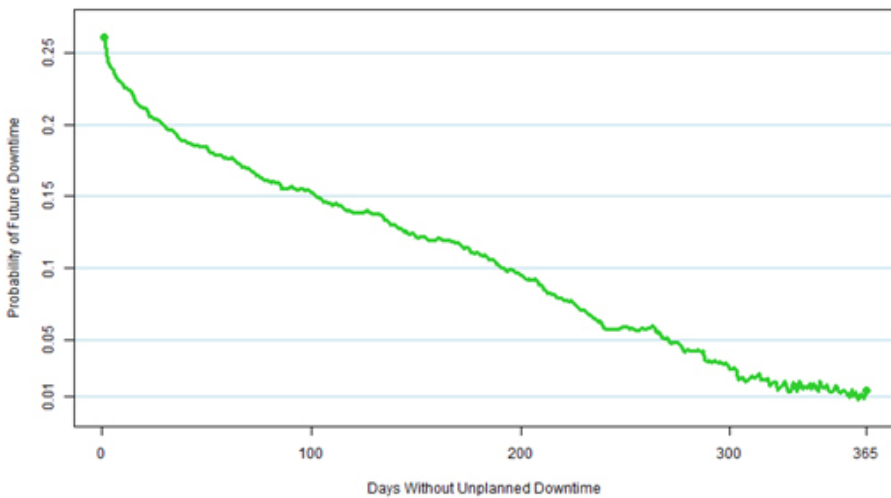


Figure 2: Probability of experiencing first unplanned downtime as a function of uptime (Days)

This effect can also be seen by viewing availability in a chronological fashion. As aggregated history increases, there is an overall positive trend in availability.

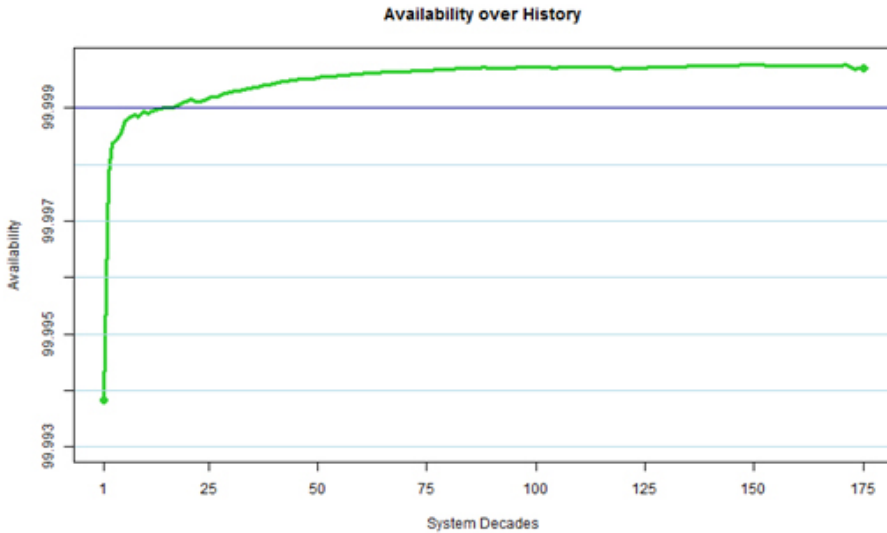


Figure 3: Availability over Nimble's history

Summary

Nimble Storage's innovative and analytics-driven approach to support and storage lifecycle management greatly increases the availability and resiliency of our storage systems. The powerful InfoSight Engine analyzes millions of telemetry data, collected daily, in order to proactively support our customer base of systems. The result has been achieving the major benchmark of greater than 'five nines' uptime:

- Nimble systems showed an overall uptime of 99.9997 percent
- As a measure of statistical confidence, problematic systems would need to occur at a rate 40 times higher than observed in order for overall availability to dip below 99.999 percent
- The InfoSight Engine continues to get "smarter" over time, being more proactive in resolving potential downtime issues. This is evident in the fact that overall Nimble system availability continues to increase over time, and the likelihood of experiencing a downtime event dramatically decreases over the course of an array's life cycle.

Next Steps

Get started with a briefing to explore how Nimble Storage can help your organization meet its availability needs. For more information about Nimble Storage, the CASL architecture, or Nimble Storage InfoSight, contact Nimble Storage by email at info@nimblestorage.com or visit www.nimblestorage.com. Also visit the NimbleConnect community (connect.nimblestorage.com) and follow us on Twitter ([@nimblestorage](https://twitter.com/nimblestorage)).



211 River Oaks Parkway, San Jose, CA 95134
Phone: 408-432-9600; 877-364-6253
Email: community@nimblestorage.com
www.nimblestorage.com



© 2012 Nimble Storage, Inc. Nimble Storage, CASL, SmartStack, InfoSight, and NimbleConnect are trademarks of Nimble Storage. All other trademarks are the property of their respective owners. WP-FNS-0214