

PROVIDING RELIABILITY WITH

GEOGRAPHIC REDUNDANCY



TABLE OF CONTENTS

| | |
|-----------|--|
| 3 | Introduction |
| 4 | Understanding High Availability |
| 6 | Calculating Availability |
| 7 | Revenue Protection |
| 8 | n-Share |
| 9 | Key Benefits |
| 10 | Ensuring Maximum Uptime |

INTRODUCTION

Commercial efforts to use IP to transmit voice calls began in the early '90s. There were numerous technical difficulties that had to be sorted and, as a result, the adoption of VoIP as a major player in telephony markets was limited. However, each improvement increased acceptability and today VoIP has become as ubiquitous as switched-circuit telephones.

A number of factors are driving the rapid growth of IP-based voice networks including cost considerations and the opportunity for businesses to increase productivity through easy information sharing, maximized utility of available bandwidth, and optimized converged network services.

As service providers continue to offer rich IP telephony services to subscribers, the need for high availability (HA) is becoming critical. HA is not a precise technology, but a goal based on specific business needs. HA improves reliability by reducing downtime, minimizing outages, and limiting the impact of outages if they do occur.

Accordingly, redundancy is indispensable for achieving HA and providing uninterrupted 24/7/365 service. In this report, we examine the importance of HA and redundancy in today's IP-enabled voice networks while describing how BCN's n-Share empowers end users to withstand failovers and maximize uptime.

UNDERSTANDING HIGH AVAILABILITY

"HIGH AVAILABILITY" refers to the ability to communicate with anyone, anywhere, at any time. For businesses relying on IP-based communications, the absence of HA can result in decreased productivity, lost revenue, customer dissatisfaction, and weakened market position.



Like other real-time applications, IP-based voice applications are extremely bandwidth and delay-sensitive. For IP-enabled voice networks to be a realistic replacement for standard public-switched telephone network (PSTN) telephony services, customers must receive the same high level of voice quality they can obtain with basic telephone services.

For voice transmissions to be intelligible to the receiver, this means voice packets must be guaranteed certain compensating bandwidth, latency, and jitter requirements to avoid dropped calls or excessive audio delays (otherwise known as jitter). HA ensures that VoIP voice packets receive the preferential treatment they require. To achieve HA, users must be able to place and receive calls during peak-load call rates, planned device maintenance, or unexpected failures.

There are two key elements that contribute to availability in a VoIP network:

1| CAPACITY

a measurement of the volume of traffic a network is engineered to handle. Typically, voice networks are engineered to handle a target peak-load capacity measured in calls per second (CPS). Target peak-load capacities are specific to each business and industry, and are based on measured busy-hour call rates. Traffic between Thanksgiving and Christmas, for example, might be the target peak-load capacity for a mail-order company's IP voice network.

2| REDUNDANCY

measures the extra capacity to be used only in the event of an equipment failure. If a primary node in the network fails, suffers a distributed denial-of-service (DDoS) attack, or is taken offline for maintenance, a redundant secondary node can take over the processing of the voice traffic.



CALCULATING AVAILABILITY

Reliability, resiliency, and availability are sometimes used interchangeably when discussing HA. While all three terms are related to the concept of HA, there are significant differences:

RELIABILITY is the probability that a system will not fail during a specified period of time.

RESILIENCY is the ability of a system to recover to normal operation after a failure or an outage.

AVAILABILITY is the ratio of time that a service is available to total time. Availability can be expressed as a mean time between failure (MTBF) and mean time to repair (MTTR).

The following table shows ways in which availability can be expressed:

| Availability | Downtime per Year | Types of Systems |
|--------------|------------------------|---|
| 90.0000% | 36 days, 12 hours | N/A |
| 99.0000% | 87 hours, 36 minutes | General PCs, fax machines, and printers |
| 99.9000% | 8 hours, 46 minutes | ISPs and noncritical business systems |
| 99.9900% | 52 minutes, 33 seconds | Data centers |
| 99.9990% | 5 minutes, 15 seconds | Redundant storage systems |
| 99.9999% | 31.5 seconds | Military defense systems |

MAKING OUTAGE REPORTING OUTRAGEOUS

In February 2012, the United States' FCC made new rules that require VoIP providers to report 30-minute outages that affects 911. Simply doing routine upgrades could create that kind of outage in many networks.

Source: "Making VoIP Geo-Redundancy Actually Work Well" by Mark R Lindsey

NO TIME FOR DOWNTIME

One in four companies has experienced a network disaster lasting 8 hours. Twenty-four percent had outage times over 24 hours. Sixty-four percent do not have sufficient disaster recovery plans.

Source: Infonetics Research, Comdisco/Bell South/Oracle Vulnerability, Contingency Planning Research

REVENUE PROTECTION

As solutions designed to address reliability, resiliency, and availability have no direct link to revenue generation, service providers often find it difficult to justify investments in this area. By quantifying the real cost of IT-induced downtime and evaluating the business-continuity options available, providers can avoid the risk of downtime without spending more on protection solutions than what the business stands to lose.

According to a report by Aberdeen Group, the average cost of an hour of downtime for IP-enabled voice networks is \$138,000, and that figure is expected to rise as an increasing variety of business functions adopt IP telephony.

The High Cost of Downtime

The average corporation has the potential to lose an average of \$7.8 million per year in downtime. Network downtime among large companies cost an average of \$32.5 million in lost productivity and revenue.

Source: Infonetics Research, Comdisco/Bell South/Oracle Vulnerability, Contingency Planning Research

The following table compares the cost of downtime for businesses with varying levels of preparedness:

| Yearly Cost Matrix | Best-In-Class | Industry Average | Laggards |
|--------------------------------|---------------|------------------|-------------|
| Business-interruption events | .3 | 2.3 | 4.4 |
| Time per business interruption | .1 | .1 | 9 |
| Total disruptions (hours) | .03 | 2.3 | 39.6 |
| Average cost per disruption | \$101,600 | \$181,770 | \$99,150 |
| Total cost of business | \$3,048 | \$418,071 | \$3,926,340 |

To determine the impact of downtime on your business and users, it is important to quantify the effect that each hour of downtime will have on profitability, productivity, and customer satisfaction.

Even a short disruption can have significant impact on retail outlets and other businesses that operate in an increasingly cashless economy.

n-Share

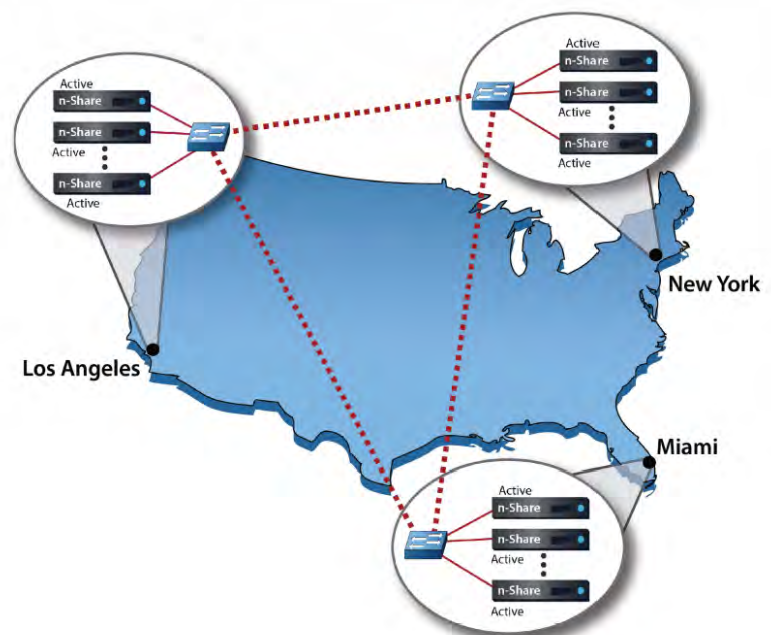
n-Share technology addresses the requirements of reliability, resiliency, and availability by providing a framework for empowering service continuity, system expandability, and ease of maintenance.

n-Share uses redundancy to ensure that call-processing systems are sufficiently reliable and scalable to handle the required number of users and devices, and resilient enough to handle various network and application outages or failures.

n-Share runs on separate general-purpose Linux machines (or nodes) created and managed specifically for, or dedicated to, the n-Share application. n-Share is deployed as an active-active redundant application, in which all nodes are online and pass call information, feature codes, voice mail, and other traffic simultaneously under normal conditions. If one node fails or loses connectivity, traffic fails over to the remaining n-Share nodes seamlessly, with no down time or loss of network connections.

n-Share is engineered to scale in a virtually limitless fashion while being able to be managed as a single application. It can be deployed either as a 2-node configuration or in multi-node clusters. n-Share nodes support geographically dispersed multi box redundancy, making them ideal for disaster-recovery scenarios

A data center with locations on the east and west coasts, for example, might install an n-Share node at each location to spread active communications across datacenters and reduce impact when disaster strikes one of the datacenters. Global enterprises, on the other hand, might deploy n-Share nodes around the world, per region or lines of business (for example, one for North America, a second for APAC, and a third for EMEA). This model enables an enterprise to address HA and redundancy across differing territorial boundaries.



KEY BENEFITS

HIGH REDUNDANCY: active/active operation enables each n-Share node to provide redundancy to all other nodes. In a hypothetical situation where all nodes in a cluster go down except one, the traffic is still processed (up to the capacity of the single node), thereby indicating the high level of redundancy available through n-Share.

EFFECTIVE RESOURCE UTILIZATION: active/active clustering actively forwards traffic to achieve maximum performance gain.

LOAD SHARING: in an n-Share deployment, the call processing load can be shared among multiple nodes for faster response times.

FAST STATEFUL FAILOVER: failover time is the quintessential index by which to judge the quality of an HA system. The faster the failover, the less impact there will be to business continuity. With n-Share, failover is as seamless as possible in order to maintain service and deliver the required user experience.

SCALABILITY: n-Share can scale vertically with capacity expansion and horizontally with performance gain.

NO SINGLE POINT OF FAILURE: in an n-Share topology, at least one alternative path for traffic flow is always available, providing the highest levels of continuity and availability.

PERFORM MAINTENANCE WITH ZERO DOWNTIME: upgrading and patching hardware and software means going offline for a period of time. With n-Share, system maintenance can be performed any time it is required, without having to plan for downtime. Minimizing downtime dramatically improves business continuity, without burdening IT, impacting performance, or affecting user productivity.

DISASTER RECOVERY: n-Share protects against site-level failures, from natural disasters to man-made activities, such as configuration errors and overwhelming service demand. You can even test your disaster-recovery scenarios in a non-disruptive manner.

PROTECTION FROM DDOS ATTACKS: n-Share's ability to replicate services across nodes offers highly effective, fault-tolerant protection against high-volume attacks designed to bring down critical business applications and resources.

Ensuring Maximum Uptime

VoIP is here to stay. When comparing investments in legacy telephone systems with those in VoIP services, it is clear that VoIP is overtaking legacy as the go-to technology of the future.

The growing popularity of VoIP and cloud-based UC has placed greater burdens on service providers to ensure maximum uptime. If a network is not available for any reason- whether due to internal issues such as operator error or infrastructure failures, or generated externally from events such as hurricanes, fires, floods, or malicious attacks- VoIP calls are not possible.

A wide-sweeping outage caused by a disastrous event can threaten the ability of you and your subscribers to continue operating and generating revenue for hours, weeks, or even days. At the other end of the spectrum, short, intermittent outages can undermine customer confidence and loyalty.

Against this backdrop, the integration of business continuity and operational availability is becoming mandatory for service providers seeking to meet and exceed the diverse requirements of their end users. With this in mind, with BCN you get a reliable, resilient, and available voice network in line with your requirements.

At its most basic level, n-Share's fully redundant, HA-driven resiliency effectively mitigates human error and infrastructure failure to ensure maximum uptime. As with most risk minimization strategies, diversification is highly effective. With n-Share, diversification is achieved using an active/active infrastructure that separates n-Share nodes from each other, so that an event that impacts a node on one network is unlikely to interrupt the nodes on other networks.