NORTEL

White Paper

What's all the fuss about network resiliency?

Network resiliency

Why should I care about resiliency?

Network uptime is critical for an enterprise striving to run in the black. Unplanned network outages can contribute to lost sales, increased overtime, loss of employee productivity and declines in customer loyalty.

Redundancy and resiliency are not the same thing. Redundant networks often have two of every network element — one device is in use while the second sits in reserve or on a shelf. Resilient networks are comprised of network devices that provide reliable failover mechanisms — either within the device or by working in concert with other network elements so that all network devices can be utilized simultaneously. A redundant network is not always the most resilient. Redundant network elements can increase network complexity and can be expensive to implement. A truly resilient network provides the maximum amount of network uptime without requiring an entire duplicate network. The exponential growth in network traffic that includes converged applications highlights the need for network resiliency.

How is resiliency measured?

Network availability or resiliency is typically measured against the standard of being 100 percent operational or never failing. A goal for many network solutions is 99.999 percent availability or "five 9s". This availability or uptime is the time during which a network or portion of the network is continuously operating or running. How much downtime is acceptable depends on the applications and revenue potential of the traffic running across the network.

Availability	Yearly downtime
90%	876 hours
99%	87 hours, 36 minutes
99.9%	8 hours, 45 minutes
99.99%	52 minutes, 33 seconds
99.999%	5 minutes, 15 seconds

Five 9s network availability may be required throughout the network either end-to-end or only across shared network resources such as the network core or backbone. The types of backup devices, availability of continuous power, N+1 redundant systems, physical separation of systems and administration practices all contribute to obtaining a five 9s resilient network.

What affects network resiliency?

Network resiliency is the sum of a number of factors. Building a resilient network involves:

- **Decreasing network complexity**
- **Eliminating single points of failure**
- **Determining the number and types of connections used and the recovery methods employed**

The complexity of a network can dramatically affect its resiliency. Complex physical and logical configurations can make troubleshooting problems very difficult. Network designs that incorporate complete redundancy of switches and routers are expensive and needlessly increase the complexity of maintaining the network. Large numbers of connections and hardware sit dormant waiting for a failure to occur. Network complexity, whether in large numbers of devices or complex configurations and topologies, can decrease the time to resolution for problems and can increase the time to implementation of network solutions.

Eliminating single points of failure in a network helps to ensure that the failure of one device doesn't bring the entire enterprise network to a halt. Designing networks without bottlenecks or choke points means that network traffic has more than one active path to a destination. Routing and switching protocols provide the first line of defense against outages. However, these protocols are only as good as the network design. Slow network re-convergence times and deterministic algorithms may not provide the network failover times required by converged applications. Out-of-sequence packets and network delays greatly affect voice, video and collaborative communications. Unpredictable delay and jitter easily create disruption in predictable network service.

Determining the connections and recovery methods used requires a thorough understanding of the network topology. How are the switches connected? How many links and how much bandwidth resides between the wiring closets and the network core? Are there redundant connections? All of these questions represent fundamental design considerations.

There are a variety of protocols designed to help networks re-converge after network errors or failures. Different protocols address different levels of recovery. However, few were designed with the current applications such as IP Telephony, multicasting or collaboration tools in mind. Fewer still boast recovery times capable of supporting these truly converged applications. Multiple protocols should be used to address resiliency at the different layers within the network.

Designing a resilient network

= hardware resiliency + network design + networking protocols

Hardware resiliency is the ability of network hardware solutions to remain available or up during a failure. Hardware elements like core routing switches need to have redundant physical attributes like N+1 power supplies as well as hot swappable cards and fan trays. The ability to swap these basic types of items on-the-fly is critical for maintaining the most basic level of network resiliency.

Network design provides the next step in ensuring a network remains resilient. The complexity of a network design contributes positively or negatively to its resiliency. Too many redundant connections and network elements can create a solution that can be difficult to troubleshoot and maintain. Too few connections or network elements may create single points of failure or traffic bottlenecks. For most enterprise solutions, five 9s reliability in the network core is attainable. Ideally, users should connect to network access points that have bandwidth-rich resilient connections to the network core. As user traffic grows, the network backbone should be able to scale as well. This makes the network core a prime target area for providing five 9s reliability. By providing the highest degree of reliability in the network core, all users can benefit from the investment. Bandwidth-stringent applications can be used network-wide based on the fact that the traditional aggregation point for network traffic — the network core — is built to support disparate types of traffic and its special needs.

Networking protocols provide a logical approach to ensuring network resiliency. Routing and switching protocols allow traffic to move around failed hardware elements or disconnected networks. However, the process of re-converging after a failure or outage may still affect application performance depending on the protocol used. All networking protocols may not provide the level of resiliency required in terms of recovery time or capabilities. Standard protocols like Virtual Router Redundancy Protocol and the Spanning Tree Protocol provide methods of network-level recovery but often lack the timing required for ensuring application continuity for VoIP and Unified Communications solutions.

The Nortel Ethernet Routing Switch 8600 provides a network core with the tools to become a five 9s resilient network. With a fault-tolerant chassis, the Ethernet Routing Switch 8600 provides hardware resiliency utilizing hot swappable components including network and CPU cards, power supplies and fan trays.

Network design is made simple with the Nortel Ethernet Routing Switch 8600. As a chassis-based solution, the Ethernet Routing Switch 8600 can be configured to both aggregate desktop traffic as well as provide core network backbone connectivity. Technologies supported include 10/100 Ethernet, Gigabit Ethernet for both copper and fiber, 10 Gigabit Ethernet, Course Wavelength Division Multiplexing (CWDM) as well as ATM and Packet over SONET. This wide variety of connectivity options means that regardless of the networking technologies used within the network, resiliency and

Figure 1: Nortel Ethernet Routing Switch 8600 Network

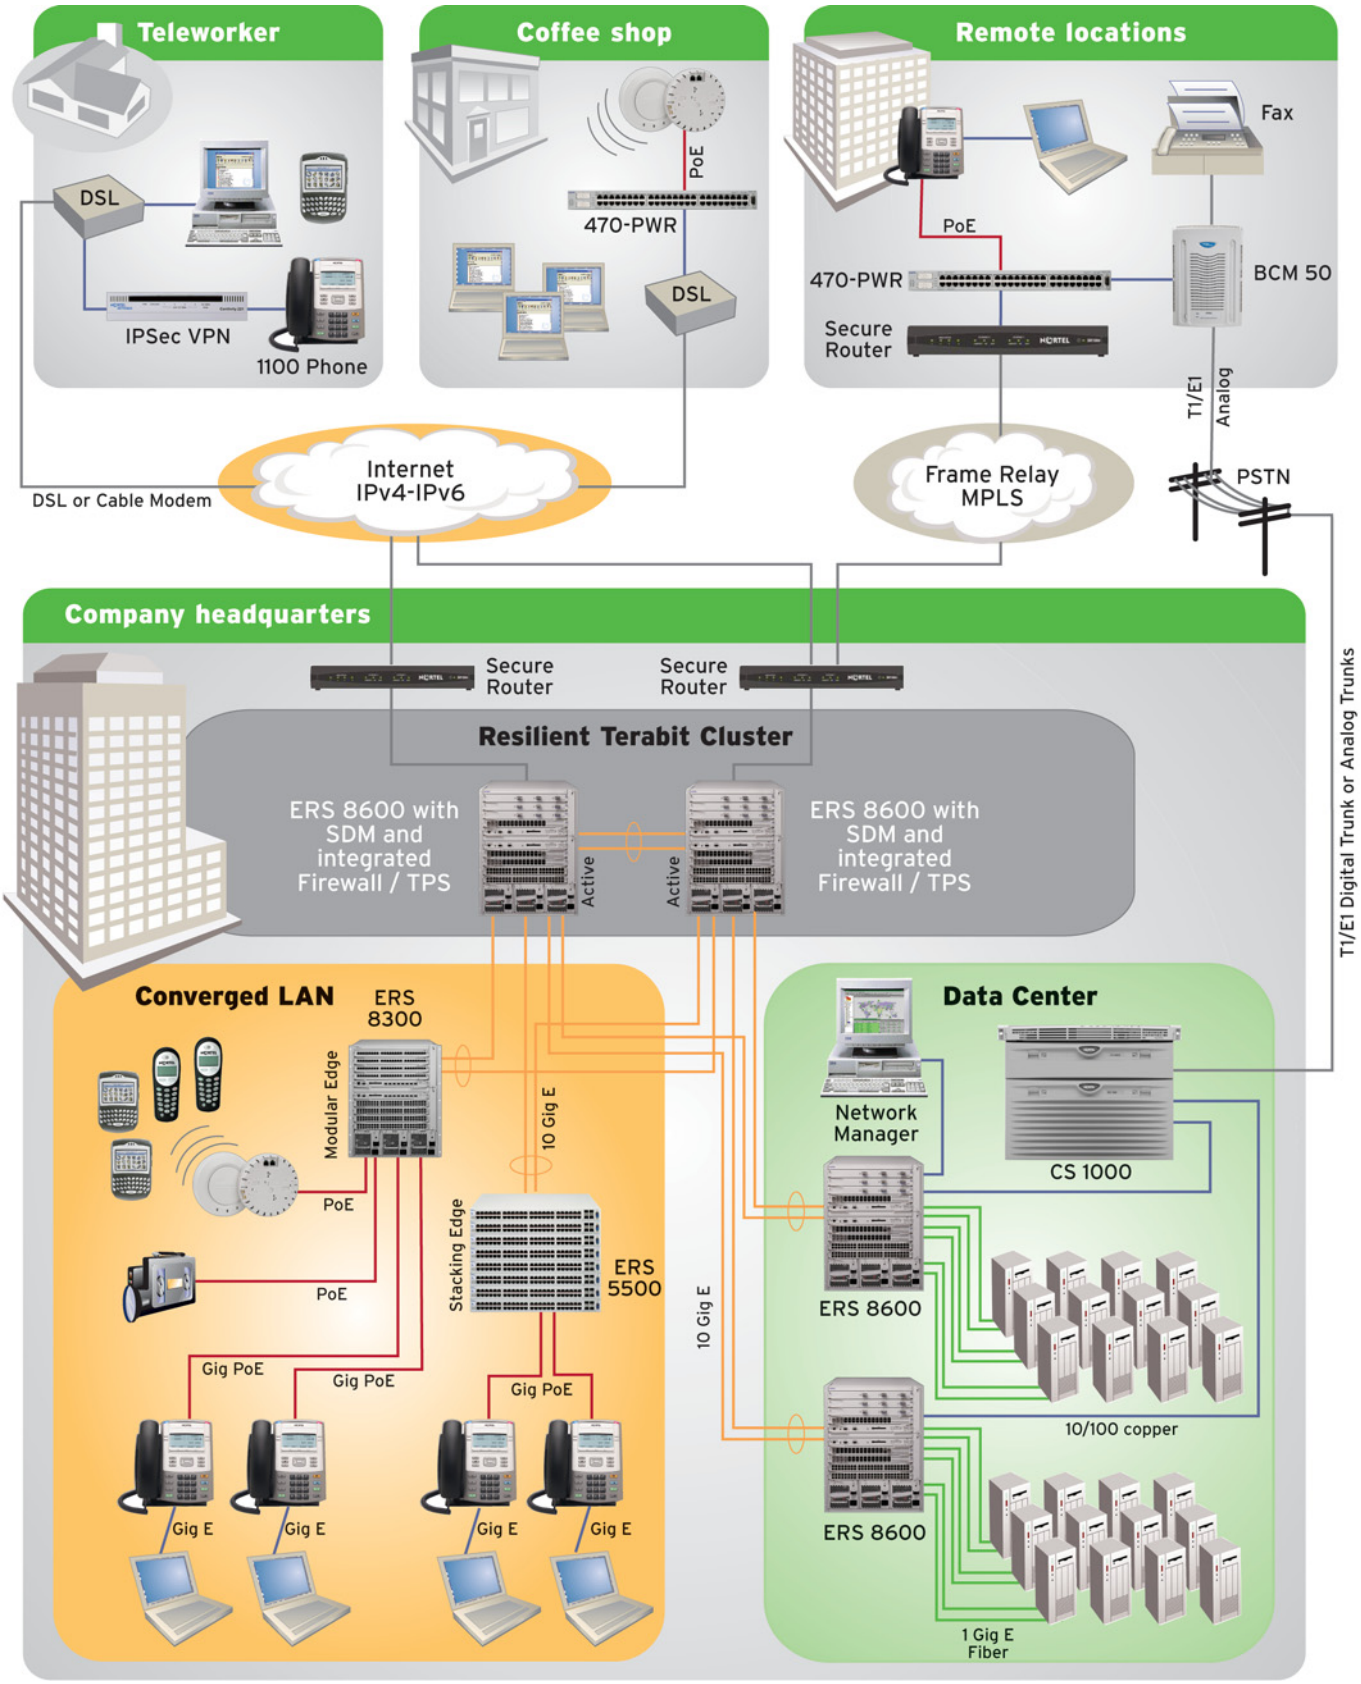

performance can be incorporated to support applications that have stringent resiliency requirements.

Network protocols for resiliency get a huge boost with the Ethernet Routing Switch 8600 thanks to Split Multi-Link Trunking (SMLT). Typical networks rely heavily on the Spanning Tree Protocol to sense trunk-level failures and switch around them. Unfortunately, re-convergence times with Spanning Tree can approach 90 seconds or more depending on the size of the network. A 90-second failure can greatly impact applications that require predictable delay and jitter in the network. Split Multi-Link Trunking allows two switches residing in the network core to logically function as one, permitting wiring closets and access points to be dual homed to a network core. This dual connectivity is unique in a number of ways. First, if one of the core switches were to experience an outage, the second switch can immediately take over the responsibilities of the first. Secondly, since the access points have dual active connections to the network core, bandwidth is immediately doubled from the wiring closets to the network core with no added expense.

Conclusion

Providing a resilient network is a multi-layered process that allows enterprises to support current applications while providing a solid foundation for their future network growth.

Working towards five 9s resiliency starts in the network core with network hardware elements, design and protocols working together to ensure reliability. By ensuring the maximum amount of resiliency in the most heavily used area of the network, network managers can feel confident that applications like IP Telephony, multicasting and collaboration tools will get the network resources when and where they are needed.

A truly resilient network provides the maximum amount of network uptime without requiring an entire duplicate network. The exponential growth in network traffic that includes converged applications highlights the need for network resiliency.

Nortel is a recognized leader in delivering communications capabilities that enhance the human experience, ignite and power global commerce, and secure and protect the world's most critical information. Our next-generation technologies, for both service providers and enterprises, span access and core networks, support multimedia and business-critical applications, and help eliminate today's barriers to efficiency, speed and performance by simplifying networks and connecting people with information. Nortel does business in more than 150 countries. For more information, visit Nortel on the Web at **www.nortel.com**.

For more information, contact your Nortel representative, or call 1-800-4 NORTEL or 1-800-466-7835 from anywhere in North America.

NN106100-101006

In the United States:
Nortel
35 Davis Drive
Research Triangle Park, NC 27709 USA

In Canada:
Nortel
195 The West Mall
Toronto, Ontario M9C 5K1 Canada

In Caribbean and Latin America:
Nortel
1500 Concorde Terrace
Sunrise, FL 33323 USA

In Europe:
Nortel
Maidenhead Office Park, Westacott Way
Maidenhead Berkshire SL6 3QH UK

In Asia:
Nortel
United Square
101 Thomson Road
Singapore 307591
Phone: (65) 6287 2877

NORTEL › BUSINESS MADE **SIMPLE**