

LinkProof:

Creating highly available, flexible,
and cost-effective multi-WAN
load-balancing solutions

January 2009

North America

Radware Inc.

575 Corporate Dr., Lobby 1
Mahwah, NJ 07430
Tel: (888) 234-5763

International

Radware Ltd.

22 Raoul Wallenberg St.
Tel Aviv 69710, Israel
Tel: 972 3 766 8666

www.radware.com



Table of Contents

Executive summary	3
WAN has become critical to business and IT	4
Datacenter consolidation.....	4
Business Continuity/Disaster Recovery readiness.....	4
Workforce mobility.....	4
Customer/supplier integration over the network and uptake of online services.....	5
Voice and video over IP.....	5
Conclusion.....	5
Challenges in designing and managing a WAN optimized for application delivery	5
Availability and Business Continuity management.....	5
Service level management (performance).....	6
Capacity planning and financial management.....	7
Optimizing application delivery WANs with multi-WAN load balancers	7
Ensuring availability and business continuity.....	7
Optimizing service levels and performance.....	8
Creating a flexible, scalable and cost-effective WAN.....	9
Gaining control with bandwidth management.....	9
Case Studies	10
Case study: Calculating the cost of downtime.....	10
Case study: Regency Hospital.....	10
Case study: Unicom.....	11
BGP4: a cumbersome and expensive solution for multi-homed networks.....	11
The Radware advantage: what to look for in a good multi-WAN load-balancing solution	12



Executive summary

Over the past few years, businesses have come to rely significantly on information technology and networked applications for day-to-day operations as well as for creating competitive edges.

IT organizations are adapting by increasing their focus on the service-delivery aspects of their responsibilities; and the role that wide area networks (WANs) play in the service-delivery infrastructure has grown considerably. This is a result of the convergence of a number of trends, namely: data-center consolidation and centralized application roll-outs (for example, desktop virtualization), increased workforce mobility, Business Continuity and Disaster Recovery (DR) awareness, customer and business-partner integration and uptake of Internet-based online services, and Voice/Video over IP (VoIP).

Radware's LinkProof provides the necessary functionality to address application down-time, latency, and bandwidth constraints, and contribute to a complete, cost-effective, and scalable WAN optimization solution for application delivery.

By routing traffic over multiple ISP links (multi-homing), LinkProof not only enables the creation of a redundant WAN architecture that addresses application up-time concerns, it also directly impacts the performance of applications and improves the speed at which they respond. This improvement is greatly enhanced by techniques such as LinkProof's Proximity™ detection and advanced health monitoring capabilities. Integrated IPS and DoS security ensure that these performance gains will not be diminished even in adverse situations such as high-volume virus attacks.

WAN scalability and cost-reductions are realized through complete freedom of link choices that LinkProof allows architects to consider (including broadband), the ease in which pipes and service providers are introduced or removed, and the more effective utilization of existing connectivity resources through load-balancing and bandwidth management techniques. This solution far outstrips BGP4's ability to provide dynamic, bi-directional link load-balancing, while doing so at a fraction of the cost and complexity; it can also provide the missing functionality to environments already utilizing BGP4.



WAN has become critical to business and IT

Over the past few years, businesses have come to rely significantly on information technology and networked applications for day-to-day operations as well as for creating competitive edges.

This includes interactive applications such as ERP, CRM, e-mail, design/development systems, employee self-service portals and hosted Web sites; remote desktop applications; real-time communication applications such as voice and video over IP (VoIP); and basic network services (typically “hidden” from the users) such as DNS, DHCP, and print services.

As a result, IT organizations are increasingly focused on the service-delivery aspects of their responsibilities, which include* service level management, availability management, capacity management, financial management, and IT service continuity management. Of all the various service-delivery infrastructure components, the importance of the role that wide area networks (WANs) play, whether private or public (for example, the Internet), is growing considerably. This is a result of the convergence of a number of trends, namely, data-center consolidation and centralized application roll-outs (for example, desktop virtualization), increased workforce mobility and Business Continuity and Disaster Recovery (DR) awareness, customer and business-partner integration, and uptake of Internet-based online services and Voice/Video over IP (VoIP).

Datacenter consolidation

The operational costs and complexity of managing multiple “mini-datacenters” at satellite and branch offices are driving enterprises across the board to remove server resources from remote locations and consolidate them in a few, highly centralized sites. Easier compliance with regulations and IT policy frameworks such as Sarbanes-Oxley, Basel II, PCI, ITIL® and COBIT, and meeting security needs such as access-control are also strong drivers. New application roll-outs are embracing the centralized model from the get-go, and a larger portion of the corporate work-force is accessing critical applications over the WAN.

An example of this trend is the move to desktop virtualization, where desktops in branch offices run thin clients remotely connected to their corresponding desktops instances running in the data center.

Business Continuity/Disaster Recovery readiness

As more and more enterprises take their Business Continuity and Disaster Recovery (BC/DR) plans seriously, the need for “anywhere-access” and “always-up” business applications (for example, verifying that remote desktop cash register applications are always accessible) has become an important part of their IT and business strategy.

Workforce mobility

An increasingly larger portion of the modern enterprises’ employees work from home or are constantly on the road. This drives a similar need to the one described above: “anywhere-access” to business applications—for example, Blackberry access from an outside sales force to the corporate CRM application.

* Information Technology Infrastructure Library (ITIL)



Customer/supplier integration over the network and uptake of online services

The Internet and associated technologies such as XML have made accessing customers and suppliers over the network a cost-effective proposition and have opened up the door to business applications that transcend traditional IT boundaries. E-commerce is a trivial example of accessing customers over the network. Another example is supply chain management (SCM) systems that tie into a component-supplier's database overseas. More examples are the use of Internet-based services such as hosted CRM (à la salesforce.com) or online package shipping (like UPS). In all these cases, the WAN or Internet are an integral part of the solution.

Voice and video over IP

Cost-cutting opportunities and advanced "converged" services in conjunction with a large supply of mature solutions are driving rapid demand for and deployment of IP-based voice and video communication solutions in geographically dispersed enterprises; and this is emerging as a strong driver for re-architecting the WAN. VoIP applications exhibit extreme sensitivity to even minute delays or disruptions to the communication path.

Conclusion

IT organizations must invest in the proper resources to ensure that their Wide Area Network meets their service-level, availability, capacity-planning, service-continuity, and financial-management needs.

Challenges in designing and managing a WAN optimized for application delivery

Major obstacles prevent smooth implementation of WAN-based application-delivery models.

Availability and Business Continuity management

First and foremost is the inability to inherently guarantee availability and continuity of a single WAN or Internet link. WAN connectivity is almost always outsourced to a service provider. A report by Infonetics* estimates that medium-sized businesses suffer around 28 hours of service-provider downtime every year, and that these outages cost, on average, almost \$200,000 in lost revenue and productivity over the course of a year, representing the single largest source of downtime cost (about 22% of the total). As long as mother-nature and human-nature remain a constant, communication service-provider disruptions are here to stay.

The following examples are split into four types: natural (disaster), human error, malicious intent, and business environment:

- **Natural (disaster). 9/05:** Hurricane Katrina wipes out most telecom facilities in the hardest-hit areas in the Gulf of Mexico. Some carriers can only keep data and voice services up and running in switches until generator fuel runs out.
- **Human error. 11/05.** Two fiber cuts to Cogent's network (Houston-Tampa, and Philadelphia-Washington D.C.) leave all of the ISP's customers without connectivity for four hours.
4/07. Insufficiently-tested software introduced into RIM's operational systems disrupts e-mail and application access by all of its US-based Blackberry customers for two whole days.

* Infonetics Research, "The Costs of Downtime: North American Medium Businesses 2006"



- **Malicious intent. 8/04.** High-tech burglars vandalize Fast24, a major UK-based ISP, leaving a large portion of British Telecom's customers without Internet service for over 24 hours.
- **Business environment. 10/05.** Two tier-1 providers, Cogent Communications and Level-3 fail to solve a business dispute. Level-3 cuts its peering connection with Cogent, resulting in the inability of customers of one provider to connect to customers of the other provider.
2/08. Tele2.fr, an ISP in France was down for days due to a network problem within the ISP's datacenter.




Figure 1 - Centralized application delivery increases the risks associated with WAN failure

Service level management (performance)

The main causes of the performance degradation that server-based applications incur when accessed over a WAN and that users experience when accessing the Internet are:

- **Bandwidth constraints.** Media rich, bandwidth-hungry applications that performed fine over a 1-Gbps LAN are challenged to provide the same transaction speeds over a 1.45-Mbps T1 or 500-kbps DSL connection. This is especially true when the WAN gets temporarily congested, or as more applications (for example, streaming video and VoIP) get trunked over the same pipes.
- **Misuse of bandwidth.** Many organizations do not monitor and/or manage the traffic passed over their Internet links. For example, internal users running P2P applications or transferring large files over the Internet link causes the link to fill up with non-business (for example, P2P) traffic, which degrades performance for legitimate traffic and users.
- **Lack of application prioritization.** Different applications may require different amounts of constant bandwidth and steady up time. Most organizations today do not bother verifying that applications requiring 24/7 up-time are indeed guaranteed that up-time.

Attempts to remedy the situation are often stymied by a lack of visibility as to what applications are traversing the network. As a Network World article* states, "Most of what is written about WAN optimization discusses technologies such as caching, compression and protocol acceleration, and

* Network World's Wide Area Networking Newsletter, 11/07/06



their ability to make applications perform better... successful application delivery requires that IT organizations are able to identify the applications that are running on the network and are able to ensure the acceptable performance of the applications that are relevant to the business while controlling or eliminating applications [such as Internet radio or streaming video] that are not.”

Capacity planning and financial management

As more and more applications converge to IP-based delivery networks, the ability to efficiently manage existing resources, and plan for future connectivity needs becomes increasingly difficult.

Extreme over provisioning is a costly and unacceptable practice in many environments. For example, a company may purchase a fractional DS3 access line anticipating future needs, instead renting a T1 line *actually* needed for the present time.

Today’s dynamic business environment requires an IT communication infrastructure that is flexible, scalable, and that can cater to sometimes opposing needs of different applications traversing the WAN. For example, the many-to-many architecture, low latency and high-performance of MPLS might suit VoIP connectivity between offices, but could be financial overkill for branch-office connectivity to a centrally-hosted CRM application, where a VPN over a couple of DSL lines might suffice.

Creating a WAN architecture that reconciles these needs and provides a cost-effective solution for the present while incorporating “on-demand” flexibility to meet future changing needs quickly is not trivial. Being able to take advantage of new, cost-effective offerings in the local markets of specific sites and working with multiple providers can ensure competitiveness. However, the risks, and operational and financial hurdles associated with introducing new WAN links, switching providers, working with more than one provider simultaneously, or migrating from one carrier-technology to another (for example, from Frame-Relay to MPLS) can be prohibitive.

Optimizing application delivery WANs with multi-WAN load balancers

Ensuring availability and business continuity

The simplest and most effective approach to dealing with WAN and Internet reliability issues is the concept of multi-homing, also referred to as multi-WAN load balancing. A multi-homed corporate network uses more than one link and/or service provider in parallel to connect to the outside world (that is, Internet or WAN, or to interconnect between sites).

Multi-homing as a concept is not new, and many organizations today implement some form of multi-homing. However, many multi-homing implementations still leave organizations at risk of downtime. Downtime is common due to a basic premise that states that having a (usually cheaper) idle backup link is a good enough solution. This premise is a misconception since, in the case of downtime of the active link (whether because the router failed or some farther point in the WAN failed), the time it takes to switch to the back-up link may cause significant business and revenue losses. Note that the time frame in which the links are switched depends largely on the method used for the switching. For example, in the case of BGP4, it may take three minutes to switch traffic between links, during which time there might not be any access to hosted Web sites.

Organizations that understand the back-up link misconception and link switching issues integrate multi-WAN load balancers, which guarantee that link availability is monitored end-to-end and that in case of link failure, traffic is switched in micro-seconds, thus guaranteeing 24/7 up time for all applications and Internet access.



With the LinkProof product line, Radware introduced a multi-WAN load-balancing technology into the marketplace, making a more practical approach to multi-homing available.

LinkProof switches fill in all of the functional gaps left open by BGP4 and other multi-homing vendors to provide a cost-effective multi-homing solution, including:

- Immediate detection of link-failures and automatic split-second failover to an available link, so that the transition is transparent to users and applications
- Layer 4–7 end-to-end health monitoring probes that extend beyond the links' gateway routers all the way across the "cloud", guaranteeing that link state is monitored even of a remote point along the WAN path fails
- Use of varied link types (for example, Frame-relay T1 combined with ADSL broadband) to create cost-effective yet resilient networks
- Simultaneous utilization of all available links and bandwidth (that is, load-balancing), so that connectivity costs are not wasted on a "dark" or poorly utilized back-up line
- No hassles dealing with ISPs, no ISP cooperation required, and no problems supporting different IP address spaces issued by different providers

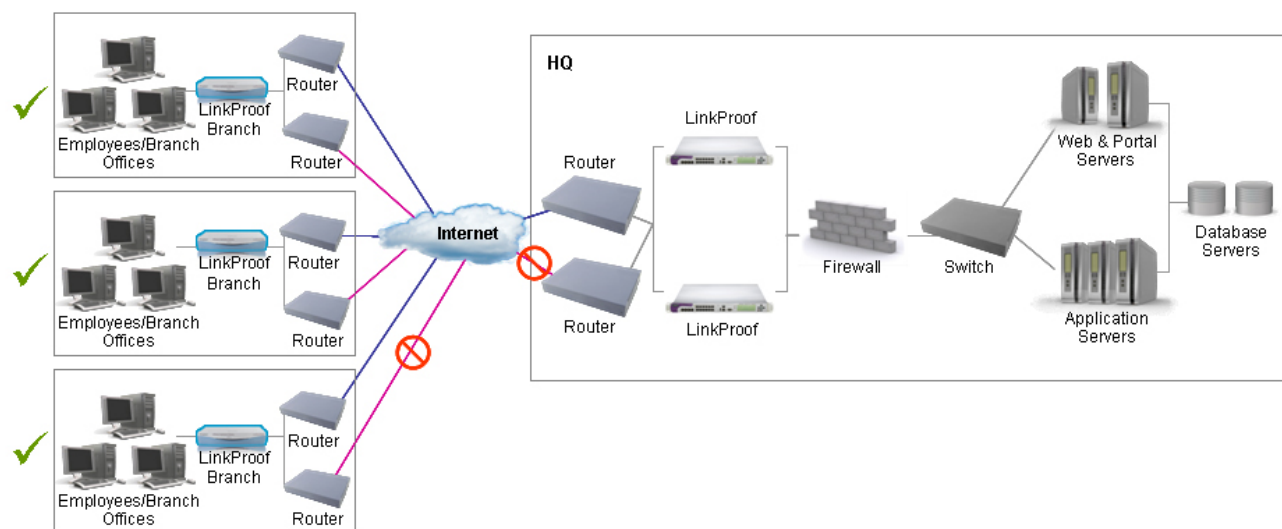


Figure 2 - High availability through a redundant WAN architecture

Optimizing service levels and performance

In addition to addressing the reliability aspect of WAN connectivity, LinkProof provides application-performance benefits and contributes to overall cost reduction of a consolidation project.

The concept of increasing the speed of application traffic by routing it through more than one network path is quite straightforward, in much the same way that adding a second lane in each direction of a highway significantly improves traffic flow despite vehicle breakdowns, construction and other congestion causes. The effectiveness of such an approach, however, depends greatly on the routing device's (the multi-WAN load balancer's) ability to quickly and accurately detect congestion, and dynamically make routing decisions to bypass roadblocks, and in a manner that is relevant to the specific type of traffic payload that is being routed. (So for example, a high-latency path is avoided for VoIP traffic, and a low-bandwidth path is avoided for bulk-file transfers). Some



ISPs might offer faster access to particular domains or destinations than others, thus finding the fastest path using Radware's patented Proximity™ algorithm, session-by-session load-balancing, and application-based routing are all examples of advanced capabilities that translate the theoretical performance improvement inherent in multi-homing into a reality.

Creating a flexible, scalable and cost-effective WAN

For those scenarios and applications where it is clear that pipe size (that is, bandwidth) will be a performance constraint, multi-WAN load balancers again lend themselves to an optimized solution architecture by utilizing the available building blocks to achieve the bandwidth goals. Instead of exclusively considering less cost-effective, single-pipe T1s or T3s, network designers benefit from flexible mix-and-match approaches of enterprise-grade and more cost-effective "business-grade broadband" (that is, cable and DSL) links, with the automatic failover capabilities of LinkProof ensuring that reliability is not compromised.

The use of multiple service providers can also contribute to making sure services are competitively priced.

With LinkProof in place at the network edges, a low-risk, cost-effective, and transparent 'add-links-as-you-grow' approach can be adopted, where additional Internet links are seamlessly added to the load balanced links pool, as the company grows its business and bandwidth needs increase. In addition, this allows the company to delay growth decisions to the optimal time.

Gaining control with bandwidth management

LinkProof's integrated bandwidth-management and traffic-shaping capabilities can even further increase the cost-effectiveness and performance of the WAN by ensuring that limited resources are allocated in sufficient amounts to critical applications first and foremost, and that non-critical uses of the WAN, for example, peer-to-peer (P2P) file-downloads, Web-browsing and other "extracurricular" uses of the Internet by employees (that can be quite bandwidth-hungry) are kept to a minimum and do not create the false sense that more bandwidth needs to be budgeted for.

The ability to make applications perform better requires that IT organizations be able to identify the traffic running on the network and ensure acceptable performance of the applications that are relevant to the business while controlling or eliminating applications that are not. Advanced bandwidth-management engines classify traffic traversing the WAN, providing administrators with a detailed breakdown of which applications and network protocols are present and what portion of available resources are being consumed by each. Granular controls are provided to reserve the optimal amount of bandwidth, limit latency for business-critical applications and ensure SLAs to customers, while "throttling down" undesirable traffic classes, for example, P2P or Internet radio.

The importance of performance visibility also extends to managing service-provider SLAs. With detailed insight into base-line WAN performance, and an ability to measure connectivity degradations and associate them with a particular provider, the task of creating manageable service-level agreements becomes possible.

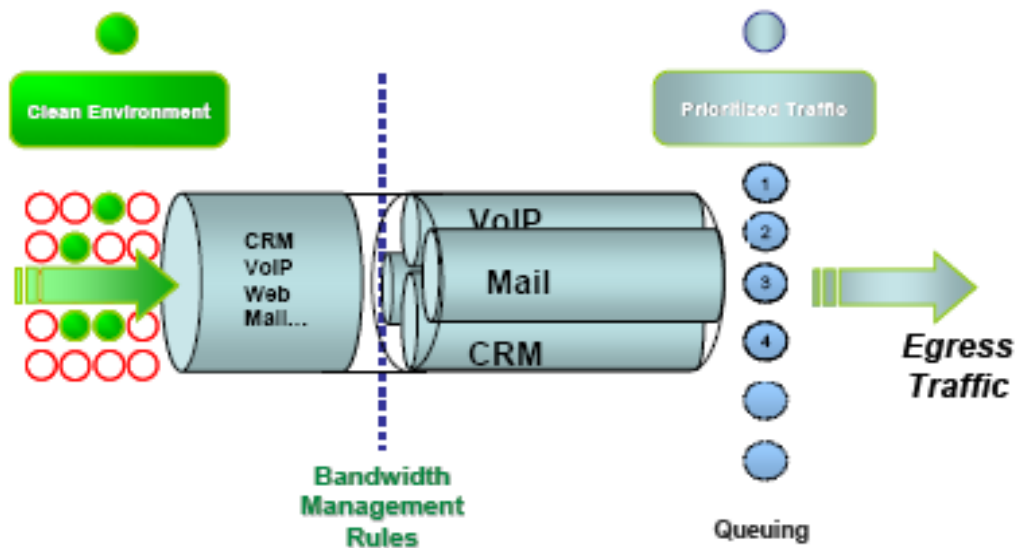


Figure 3 - Effectively classifying traffic and allocating limited bandwidth resources

Case Studies

Case study: Calculating the cost of downtime

ACME is a medium-sized retail company with 500 employees. The company sells its products via its Web portal located in the company's single datacenter, generating an average of \$50,000 an hour.

The company has a single internet link with a guaranteed SLA of 99.9%, this translates to roughly nine hours of downtime a year.

The annual loss of revenue for ACME per year, due to link failure or ISP network problems is calculated to be \$450,000 ($\$50,000 \times 9 = \$450,000$).

Add to this the price of the Internet connectivity, which must be paid regardless of link state, the annual loss of revenue for ACME per year may reach \$480,000.

Case study: Regency Hospital

Based in Alpharetta, Georgia, Regency Hospital provides intensive care to medically complex patients using aggressive clinical therapies, advanced equipment, and an interdisciplinary team approach centered on individual patient needs. The typical patient is critically ill and requires a longer length of stay in an intensive care environment than traditional short-term hospitals are designed to provide. To properly administer care, Regency's staff requires an immediate, uninterrupted, and secure VPN.

“By the time Regency opened its tenth hospital, it was clear they needed more bandwidth and redundancy at the corporate data center to ensure uptime for the hospitals – especially with the new online clinical systems that were being planned for deployment. The dual LinkProof solution was the perfect match for these requirements, allowing Regency to utilize two multi-T1 pipes from independent carriers (AT&T and Sprint).” *David Hampson, president of Enroute Networks*



Case study: Unicomm

Unicomm of Italy owns and operates over 100 supermarkets and convenience stores across the country. The company relies on continuous network connectivity to complete its business transactions. Unicomm's previous solution used a dedicated line or CDN-based solution, and all Internet connections were conducted over a virtual private network (VPN), using ISDN as backup, which was not a sufficient or scalable solution for the large amount of data Unicomm's growing store count required.

"Unicomm is rapidly opening new stores to expand its business. This places increasing demand on the network. We needed a cost-effective solution that would bolster Unicomm's existing network structure while ensuring reliable access and service to their heavily trafficked network load."

Federico Vecchiati of Miriade IT Consulting

Using LinkProof, Unicomm was able to replace its CDN with high-speed ADSL and HDSL lines, thus increasing the bandwidth available to stores. It could also maintain its existing VPN, saving LAN reconfiguration. No longer restrained by a single service provider, Unicomm was able to evaluate a variety of network providers and take advantage of special offers and market competition. As a result, network operating costs were slashed by almost 50%.

BGP4: a cumbersome and expensive solution for multi-homed networks

Border Gateway Protocol 4 (BGP4) is used by major carriers to route Internet traffic on the backbone, often referred to as "the cloud." The current version, Version 4, has been in use on the Internet since 1994. While tempting to think that BGP4's functionality can be extended to create multi-homed network topologies on corporate premises, it is a cost-prohibitive and very inflexible solution, which rarely meets the needs of typical organizations and enterprises.

BGP4 requires the use of middle-range to high-end routers, for example, Cisco 2821 XM routers with two High-Performance WAN Interface Cards (HWICs) or Cisco 7200 series with appropriate extensions, with costs running up to \$10,000 or more per router. Due to its static nature, constant administrative tweaks are necessary; and this typically requires the support of full-time, appropriately trained staff. Additional costs include obtaining an Autonomous System Number (ASN, which is not available in all countries) as well as paying monthly peering fees to the service providers.

Embracing BGP4 essentially means that the corporate network becomes part of the backbone, with ownership and manageability distributed between multiple carriers. When problems occur, finger-pointing or lack of cooperation can seriously slow down troubleshooting attempts.

BGP4 is far from a complete multi-WAN load-balancing solution. It lacks much of the control offered by more up-to-date solutions. *It offers no inbound traffic load-balancing*, and even outbound load-balancing is not very granular, happening only at the subnet-level instead of per-unique destination, and not adapting to dynamic *changes in link response times* and degradation in performance. Because BGP4 was designed to maintain stability in the cloud, it typically supports very slow path convergence. This means that in the case of a down link, failover can take 2–20 minutes instead of being immediate. Finally, BGP4 requires symmetrical, carrier-grade architectures, that is, all links must be identical and T1 or above; mix-and-match architectures that include more cost-effective links (for example, cable, DSL or wireless) are not possible.

For environments that currently employ BGP4, *LinkProof can easily be deployed in parallel* to address the functionality issues described above, such as link health-monitoring and dynamic adaptation to degrading response times.



The Radware advantage: what to look for in a good multi-WAN load-balancing solution

Radware pioneered multi-WAN load balancing with the introduction of LinkProof and has maintained market leadership through patented innovation,* the broadest product line, and the most comprehensive set of integrated capabilities designed to meet a business's WAN availability, performance, scalability, and cost needs.

The following LinkProof features comprise reasons for including LinkProof in a company's multi-WAN load-balancing shortlist:

- **Broadest WAN load-balancing offering** - LinkProof's product line offers the widest range of deployment options from 5-Mb to multi-gigabit bandwidth, designed to suit branch offices, SMBs, service providers, and large enterprises.
- **Meets the needs of all company applications** - Different businesses use Internet and WAN connectivity for multiple and often varied needs. Addressing these needs, LinkProof provides:
 - **Advanced inbound traffic load-balancing using integrated DNS routing** - supports companies that host their own Web site or self-service portal internally.
 - **Advanced persistency and application grouping features** - ensure that load-balancing will not disconnect users' interactive applications in mid-session.
 - **Split-second failover and controls to limit latency to ensure VoIP call quality** - Radware has earned top ratings in an extensive lab test of VoIP traffic resiliency and quality, performed by independent test center Miercom. For more information on the results of the test, go to <http://www.radware.com/content/document.asp?v=about&document=7500>.
- **On-demand multi-WAN link load balancing solution** - enables a company to cost-effectively upgrade its connectivity solution through flexible mixing of link types and a transparent '*add-link-as-you-grow*' approach with no extra integration costs. On-demand scalability in throughput and WAN connectivity enables adding more links, services, and platform throughput as the company grows its business.
- **Total freedom in choosing connectivity options** - any combination of link-types (for example, frame-relay T1 combined with ADSL broadband) can be used to create cost-effective yet resilient networks. ISP introduction and removal can be timed to suit the company's needs.
- **Reduced connectivity costs** - beyond the use of bandwidth management, techniques such as load-balancing and cost-based routing (that is, flexible, cost-tier-based link selection that minimizes monthly connectivity bills) allow the total usage of available bandwidth to be allocated according to business priorities, prevent waste on "dark" or poorly utilized back-up lines, and help delay bandwidth increment decisions.
- **Easy to install and troubleshoot** - easy installation and configuration (using first-time installation wizards), good documentation, and 24/7 tech-support are available through Radware and its network of resellers, coupled with robust diagnostic capabilities like CPU and traffic utilization, and "internal ping" to facilitate technical support.

* See US Patent 6,665,702 – "Load Balancing"



- **Reliable product and vendor** - switch-based equipment is superior to less reliable, PC-based appliances. LinkProof's mean time between failures (MTBF) typically extends over 700,000 hours, and extra redundancy is available through support of redundant-pair configurations. LinkProof is a field-proven solution with thousands of customers world-wide. Furthermore, Radware is a market-leading network equipment vendor with a solid record and over 10 years of switching experience.
- **Advanced health-checks for failover** - lines can appear to be alive, but in effect, be dead. Even though an ISP's closest gateway router might be accessible, routing failures further along the ISP's backbone or peering-points might render that link unusable. LinkProof utilizes fully configurable Ping, TCP, UDP, SNMP, as well as application probes to provide full-path failure detection, where any failure in the path from source to destination through one ISP renders that ISP as "out of service" and forces failover.
- **Addresses application performance degradation, not only downtime** - LinkProof not only detects failed links and immediately redirects traffic to available resources, but also measures response times, and directs load-balanced traffic to the fastest responding links. Through patented Proximity™ probes, parameters like hop-count and latency are measured per unique destination, and user requests and traffic are redirected to the service-provider offering the optimal path to and from that destination. Furthermore, the ability to perform these routing decisions based on application type further enhances the performance of specific, critical applications.
- **Offers bandwidth management and traffic-shaping** - when WAN latency and protocol inefficiencies are not the problem, lack of bandwidth typically is. LinkProof's granular control of how the limited pipe resources are allocated ensures that critical applications receive the bandwidth they need and are passed through with minimal latency.
- **Bullet-proof, integrated, application-level security provides protection against a wide-range of application security threats and DoS/DDoS attacks** - network congestion and downtime are not always the result of benign events on the WAN. Malicious attacks can bring the network to a standstill by congesting the network with high volumes of illegitimate traffic, for example, mass-mailings by worms. Dealing with such threats and removing them from traffic flows is an additional aspect of application-performance assurance over the WAN. Offering DoS, DDoS, and IPS capabilities in-the-box removes the need for additional hardware to maintain at branch and datacenter locations.
- **Integrated VPN termination** - along with the increased adoption of cost-effective, public broadband media for back-up and primary WAN connectivity, especially at branch offices, comes an increased need for VPN functionality to protect the data traversing those public lines. LinkProof's integrated IPsec VPN termination enables that functionality to be implemented without the additional overhead and costs of dedicated VPN hardware at remote locations.
- **Support for BGP4 environments** - LinkProof can be installed painlessly in environments that already use BGP4, and provide missing functionality such as link health-checks, adaptive load-balancing based on dynamic response-times (inbound and outbound), and immediate failover.

LinkProof's multi-homing and bandwidth management capabilities provide the crucial high-availability and performance enhancements necessary for a comprehensive WAN-optimization solution.