FEATURE BRIEF: FAILSAFE DELIVERY

Introduction

Today's businesses require 100% network uptime. Unfortunately, this is not always possible because many wide-area networks (WAN) fail to deliver on the resiliency requirements of critical applications. According to Gartner the average large enterprise racks up nearly 90 hours of downtime a year, resulting in nearly $4M in losses. Infonetics pegs that loss to be much higher, citing the difficulty businesses have tracking, quantifying, and responding to downtime caused by network issues.

It's fair to ask how this is possible when WAN optimization technologies and overlays have managed network traffic delivery for more than two decades. The truth is that WAN optimization and overlay networks do nothing other than react to and slow down traffic when loss or congestion occurs. Often the networks that the traffic traverses are not under the application user's control because traffic passes through other transit networks. WAN optimization techniques such as deduplication, compression, latency optimization, caching/proxying, forward error correction, protocol spoofing, traffic shaping, equalizing, and rate limiting are necessary but not enough.

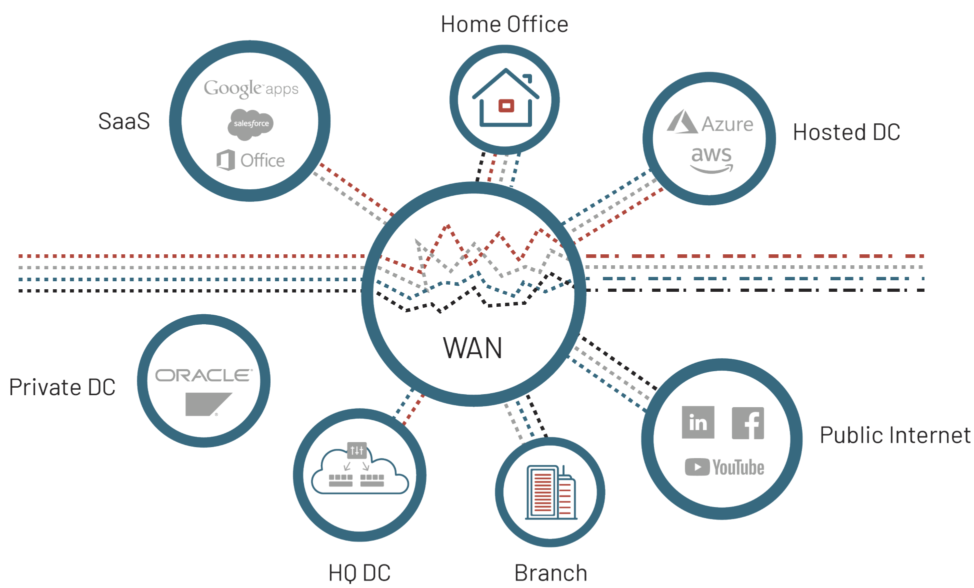

Today's SD-WAN solutions use these techniques to deliver traffic over hybrid networks, which reduces the cost of buying expensive circuits but ignore the delivery needs of modern applications. The 128T Session Smart™ Router by 128 Technology supports key WAN optimizing techniques while also providing a new integrated feature set that enables a persistent, relentless, and resolute ability to deliver traffic in the face of any and all network failures.

The 128T Session Smart Router uses a "failsafe delivery" model, routing around failures and network degradation in third-party networks, figuring out the best possible paths in real-time, utilizing uncongested servers, and delivering the SLAs needed for applications. This model ensures 100% uptime with lossless delivery.

Network Outages and Degradation

Every WAN optimization, and now SD-WAN vendor, is in business to deliver traffic. But these solutions only attempt to slow down or shape traffic when they notice degradation along the path over which they have no direct control. Even the world's top businesses suffer from downtime. It is said that Amazon's recent 20-minute outage cost the company $3.8M. This does not include the loss in opportunities, productivity, reputation, and other intangibles that cannot be quantified.

The most common causes for network outages or degradation include:

Hardware Failures: Routers today are equipped with numerous resiliency mechanisms. However even the best ones are susceptible to failures. Telstra faced a network outage due to hardware failure in March 2016. Software/hardware faults, fiber cuts, network congestion, and other errors regularly causes disruption in network services.

Configuration Errors: Despite the best intentions, numerous checks, automation, and validations – network disruptions due to configuration errors are common. Reports state that an average enterprise suffers about 90 minutes of application downtime roughly 13 times per year. Facebook suffered an hour-long outage due to a configuration error early last year.

Routing Issues: Changes in routing can affect parts of the network in ways that cannot easily be recognized. Level 3's network issues actually slowed down the entire Internet for a while in 2017. Websites such as Yahoo and eBay along with gaming services on Xbox Live were affected.

Instead of assuming that the WAN can deliver unsurpassed connectivity, the 128T Session Smart Router turns this concept on its head. Failsafe delivery assumes that the network is not perfect. It may fail. However, traffic must be delivered as long as there is any connectivity between the two end points. Unless the network is cut into two pieces with no interconnection, the 128T Session Smart Router delivers traffic to the end destination delivering lossless transport and 100% uptime.

Failsafe Delivery

The 128T Session Smart Router combines three different key technologies related to delivery of application flows to ensure optimized traffic delivery in the face of all odds. These three technology areas include:

- Intelligent Path Monitoring
- Lossless Application Delivery
- Optimized Server Heuristics

Intelligent Path Monitoring

128T Session Smart™ Routers, based on session-oriented networking paradigms, connect endpoints to services. For effective delivery of traffic loads it is important to monitor network paths between 128T Session Smart Routers in a network.

Out of Band Signaling

When no traffic is flowing between the 128T Session Smart Routers an out-of-band mechanism is needed to monitor the network path. Bidirectional Forwarding Detection (BFD) is used to gauge the health of the distributed components as well as the integrity of the network connectivity between them.

BFD [RFC5880] is traditionally used to ensure path connectivity between two routers' forwarding planes. In addition to basic BFD, the 128T Session Smart Routers use an enhanced version of BFD to measure latency, packet loss, and jitter between themselves. This data is used to supply each 128T Session Smart Router with real-time link attributes that can affect how traffic is delivered.

When a 128T Session Smart Router sends BFD packets to a peer, it includes in the BFD payload the quality point value configured for that link to its adjacency. This is retained by the receiving 128T Session Smart Router and used when making service routing decisions.

Dynamic updates to these quality points can also be done to effect changes in traffic patterns.

BFD control packets are also enhanced with “BFD metadata”. The metadata inserted by a 128T Session Smart Router can optionally be encrypted. The implementation of BFD does not change the protocol’s behavior, messages, or encoding; the only difference is the addition of 128T specific metadata when a 128T Session Smart Router transmits BFD to another 128T Session Smart Router.

Each 128T Session Smart Router uses both BFD’s asynchronous mode and echo mode. Asynchronous mode is used for liveness checks and exchanging packet loss data. Echo mode is used for determining path latency and jitter.

Connectivity

Assuming BFD connections are established between two 128T Session Smart Routers, if a router fails to receive a number of consecutive control packets from a counterpart, it treats that peer as unreachable. This has the effect of removing that router as a potential target for new session assignments.

Quality

Unique to the 128T Session Smart Router is the use of BFD for measuring link quality (latency, jitter, packet loss) between routers. It is understood that packet queuing on either the transmitting or receiving 128T Session Smart Router can skew the results of the test; if BFD packets are treated with a higher priority than session traffic, the BFD test will produce results that are more favorable than the results that the session traffic will experience.

The 128T Session Smart Router proposes to treat BFD with a relatively low priority. A 128T Session Smart Router will use echo mode BFD to periodically test each destination router; the process for doing so is for a transmitting 128T Session Smart Router to send a series of BFD packets to the destination spaced evenly. As the packets return back, the round trip time is halved to estimate the one-way latency, and the variation in inter-arrival time of the return packets estimated jitter. The number of packets sent in each series and the number received is used to determine loss.

This enhanced version of BFD running between 128T Session Smart Routers helps in network path monitoring when there are no flows.

In-Band Signaling

When traffic is flowing between 128T Session Smart Routers, the session-oriented paradigm enables the generation of detailed session data records for all flows. In addition, some in-band performance monitoring techniques further enhance this capability to enable SLA verification, troubleshooting (e.g., fault localization or fault delimitation), and network visualization.

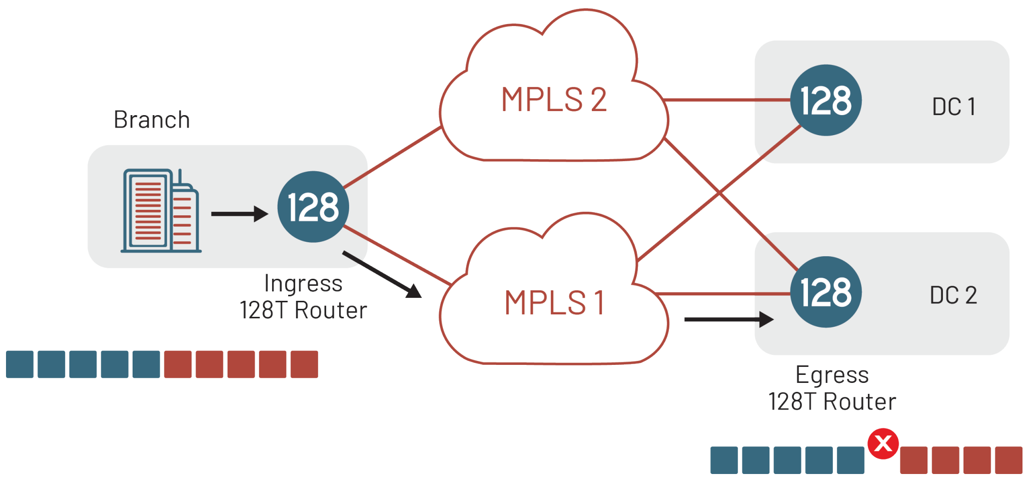

One method for flow performance monitoring is marking of packets in a specific session to different colors, thereby dividing the flow into different consecutive blocks. Packets in a block have same marking and consecutive blocks will have different markings. This enables the measuring node to count and calculate packet loss and/or delay based on each block of markers without any additional OAM packets.

Details from the egress node are communicated to the ingress node using reverse metadata or via a federation protocol. This enables in-band traffic monitoring and also enables network path monitoring when there is a flow. The data gathered from network path monitoring is used as cost metric in the algorithms used for distributing server loads. This ensures that only those servers which are reachable via paths that meet the application demands are used.

Lossless Application Delivery

The 128T Session Smart Routers can utilize innovative server load monitoring and intelligent network path monitoring to ensure that the algorithms have the best possible heuristics available to choose the most appealing paths and servers.

Redundant or alternate paths between nodes in a network can be used to reroute traffic, improve resiliency, and maximize throughput. These maximally diverse paths can provide link and node protection for 100% of paths and failures as long as the failure does not cut the network into multiple pieces.

Multi-Path Session Migration

Multi-path session migration refers to the ability to migrate an existing session to an alternate path between two 128T Session Smart™ Routers. Multi-path session migration requires that packets can be forwarded not only on the shortest-path tree but on another maximally redundant path. This guarantees 100% recovery for single failures when the paths are completely disjoint.

The administrator can configure multiple paths between two 128T Session Smart Routers. These paths can be redundant (completely disjoint) or maximally redundant (as disjoint as possible). The administrator can configure the sessions traversing these routers to work in primary/backup mode or in load balancing mode. There may be sessions which the administrator may configure as not to take the alternate path. The 128T Session Smart Router will switch traffic to the alternate path when it detects that the existing path has a failure or a link gradation that renders it unfit for use for the application.

LTE/Wired Connections

Service providers or large enterprises can provide their subscribers or branch offices with access to fixed and mobile networks. It has become desirable to use these heterogeneous networks as backups in case of failures.

In most cases the fixed wired connection is used as the primary. LTE (Long Term Evolution) or 3G connection is used as the backup. The traffic always flows over the wired connection.

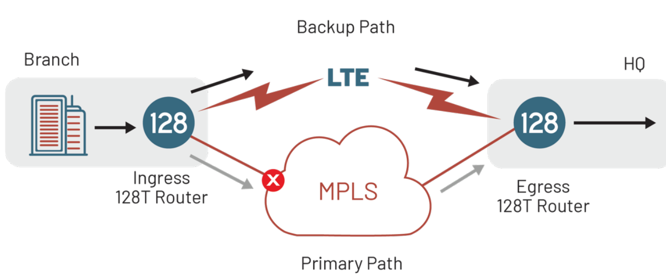

In case of failure the sessions are moved to the wireless connection. When the wired connection is restored the sessions are moved back to the wired connection.

Internet/MPLA Connections

Service providers or large enterprises can provide their subscribers or branch offices with access to connections via dedicated MPLS circuits and Internet connection. It has become desirable to use the Internet connection to save costs while the MPLS circuits act as backup if needed.

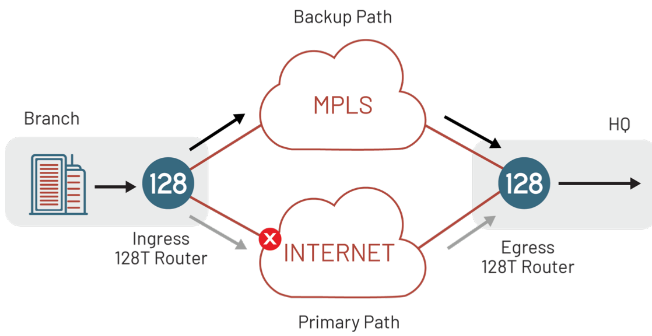

The traffic always flows over the Internet connection. In case of performance degradation, the sessions are moved to the MPLS connection. When the Internet connection is restored to acceptable performance, the sessions are moved back to the Internet connection.

Wired/Wired Connections

Service providers or large enterprises can provide their subscribers or branch offices with access to two different fixed connections for diversity. It has become desirable to use both these networks and they act as backups of each other in case of failures. It is also possible to have two paths within the same network. In this case it is desirable that both paths be used in load balancing mode.

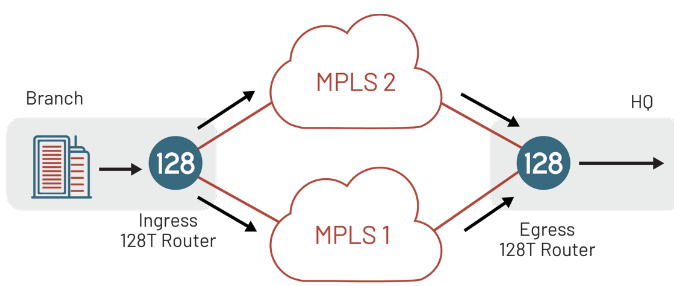

Multi-Path Session Redundancy

Packet loss is unavoidable in large networks. This loss can be due to congestion; it might also be a result of an unplanned outage caused by a flapping link, a link or interface failure, a software bug, or a maintenance person accidentally cutting the wrong fiber. Since UDP/IP flows do not provide any means for detecting loss and retransmitting packets, it is left up to the higher layer and the applications to detect, and recover from, packet loss.

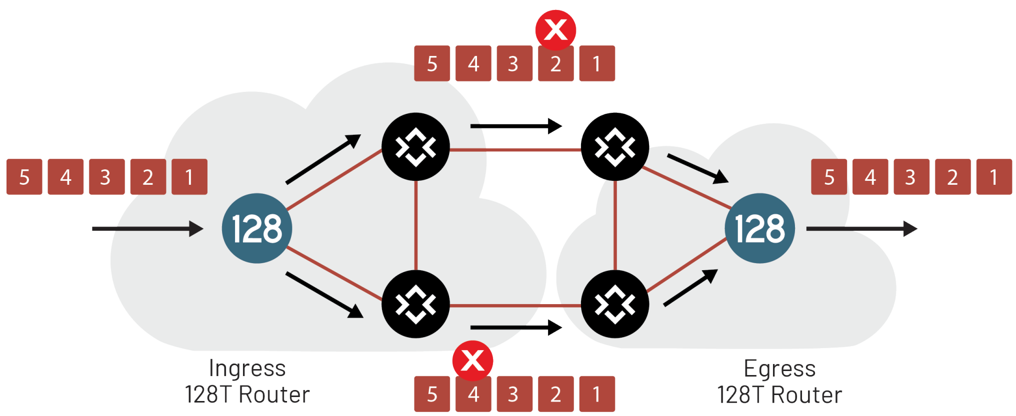

Existing loss mitigation techniques such as retransmission, forward error correction (FEC) – both media independent and media specific FEC, and interleaving have proven to be successful in limited scenarios only. One technique to recover from packet loss without incurring unbounded delay is to duplicate the packets and send them in separate redundant streams. The probability that two copies of the same packet are lost in quite small. This scheme has comparatively high overhead in terms of bandwidth as everything is sent twice. The recommendation will be to use this scheme for high-value traffic.

For example, a service provider may use this scheme to provide a superior user experience for telepresence traffic for VIP customers. The administrator can configure dual paths for a high-value session. The ingress router will duplicate the packets arriving from the session and send it over the dual paths. Both streams carry the same payload with identical sequence numbers. This allows the egress router to identify and suppress the duplicate packets, and subsequently produce a loss-free and duplicate-free output stream.

This also reduces the delay when packet loss occurs. An unrecoverable loss happens only when two network failures happen in such a way that the same packet is affected on both paths.

This technique requires the forwarding delay of the network paths to be more or less the same to ensure that the removal of duplicates and the application succeed. The egress router monitors the delays over the dual paths and reports whether they are acceptable for this scheme to function. If the delays are above the acceptable limits, the ingress router will stop duplicating the stream after a wait period until the delays are below the acceptable limit before resuming the duplicate the stream. This technique ensures lossless transport of traffic.

Multi-Path Session Maximization

When multiple paths are available between peer routers, it may be desirable to combine these networks together to offer increased access capacity.

For example, a service provider may combine speeds of a low-bandwidth wired and LTE connection, or multiple wired connections to offer increased speeds. It will also enable resolution of the mammoth flow problem. The flows are always filled in the order of how they are configured by the operator which defines the priority. This is done to ensure that a low cost medium such as wired connection is used before a wireless connection.

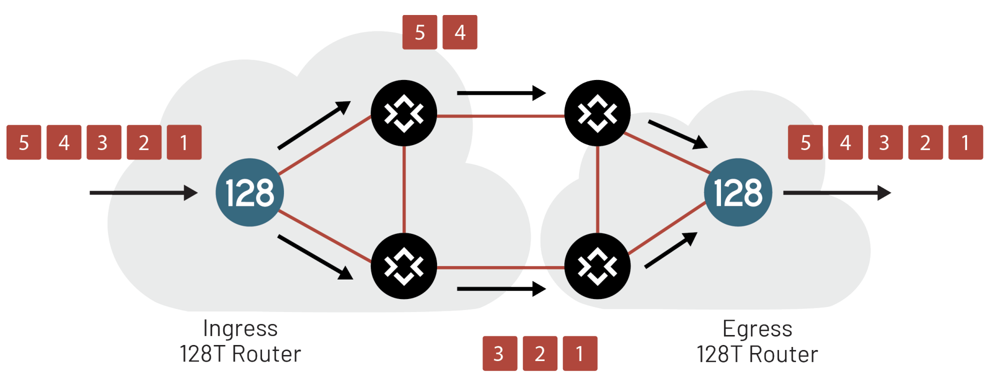

When traffic volume exceeds the bandwidth of the primary connection, the excess is offloaded to the next alternate path. The ingress router is configured by the administrator to use the primary path. When traffic exceeds the bandwidth of the primary connection, the excess traffic is offloaded to the next alternate connection. The administrator can define the interfaces that are being combined for the purpose of session maximization.

Any session using these interfaces can benefit from the combined bandwidth. The administrator can also mark flows that follow and do not follow this scheme. The scheme is dependent on the difference in transmission delays of the combined connections. If this exceeds a given threshold for a certain period, then the ingress router stops offloading traffic to alternate paths. It will resume offloading when the transmission delay has fallen below the acceptable threshold. These techniques ensure the best use of network resources and the lossless delivery of application traffic for applications.

Optimized Service Heuristics

The 128T Session Smart™ Routers operate on the notion of sessions, which are targeted to deliver traffic to service agents representing application servers. These servers are assigned loads and quality points.

Server Loads

Server loads define the ratio of load that a particular server can take, compared to other servers for the same application. The load can specify the maximum number of sessions that a server can support or the maximum bandwidth of the traffic that can be sent to the server. The load balancing algorithm of the 128T Session Smart Router will take this into account while distributing application flows across the servers.

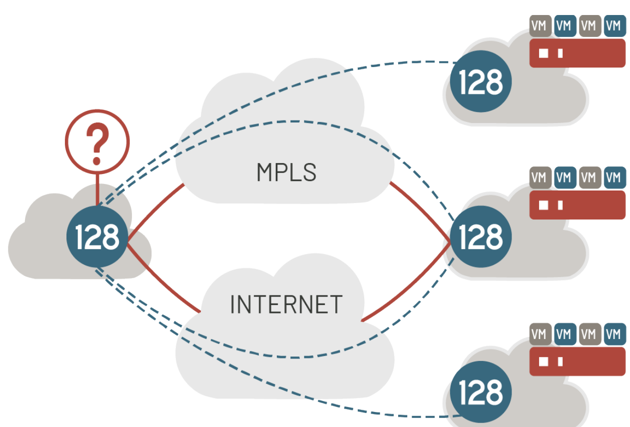

For example, consider a scenario where there are three servers which can serve a particular application and the loads assigned to them are 100, 50, and 50. A proportional distribution algorithm would spread the application flows in the ratio of 2:1:1 among them, in the absence of other factors. The 128T Session Smart Routers monitor the loads in use to a particular server and these are used for spreading loads to ensure that the servers are not overwhelmed.

Quality Points

Quality points are administrator assigned cost metrics to a 128T Session Smart Router adjacency that indicates path preference. This ensures that traffic flows to locations using links that meet SLA thresholds.

For example, if two server locations are reachable have quality points of 200 and 100 then the server with the quality point of 200 will be preferred over quality point 100, in the absence of other factors. In addition to quality points, other real-time cost criteria such as maximum session rate, packet loss, latency, and jitter can also be used to determine if the path to a particular server is suitable for use by the application. An administrator can ensure that traffic flows to the preferred server over the preferred link by assigning appropriate server loads and quality points. These metrics will change dynamically during the course of operation depending on traffic flows, leading to server loads being distributed over different WAN links.

Strategies

The 128T Session Smart Router can distribute traffic based on different criteria in two different ways. These algorithms attempt to maximize the loads to agents based on available capacities while minimizing the cost of the paths to the selected agents or servers.

Proportional

The Proportional algorithm distributes sessions to agents weighted on the relative available capacity. For example, if an agent has double the amount of available capacity as another, it will receive twice the number of sessions. This is good for distributing loads over all paths that meet the application SLA and where there is no restriction on the path being used. For example, sending all traffic to multiple servers over an MPLS network.

Hunt

The Hunt algorithm routes all sessions sequentially to the agent with the highest capacity until a load threshold is reached for the agent, it then switches to the agent with the next highest capacity, and repeats. For agents with the same capacity, the agent's name is used as a tiebreaker, lexicographically lowest first. This ensures maximum use of lower cost links because the higher cost link will only be used when the lower cost link usage has reached a threshold. This is a key strength in the SD-WAN use case which is designed to leverage lower cost links before higher cost links.

Service Uptime Check

A server uptime check is needed to ensure whether the resource is available when no traffic is flowing to the server. 128T Session Smart Routers can monitor path qualities to other 128T Session Smart Routers using routing protocols and innovative heuristics including load monitoring, however servers do not support the full routing stack to participate in this.

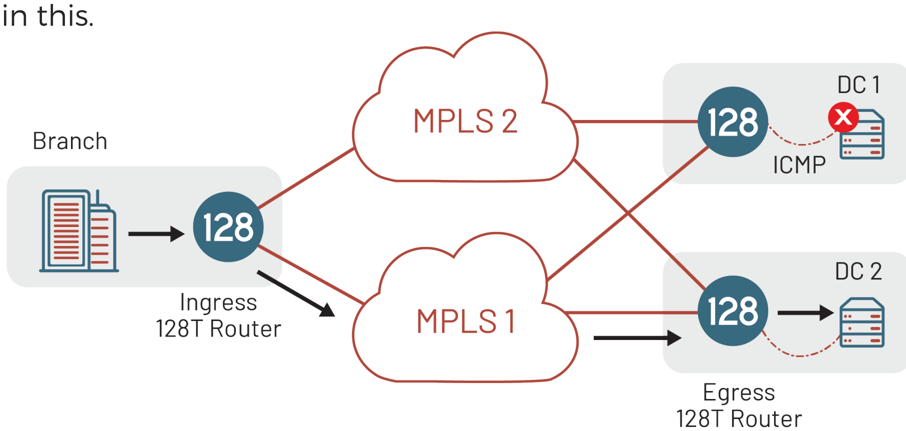

A last mile health check is required to ensure that the server is available. When traffic is flowing, the 128T Session Smart Router can use L4 and L7 techniques to determine server health. When no traffic is flowing, the 128T Session Smart Router uses ICMP to monitor last mile connectivity. This provides a quick method to verify connectivity and spread that information to ensure that 128T Session Smart Routers making decisions to distribute server loads do not include a nonresponsive server in the decision making algorithm. These mechanisms ensure that the best possible or preferred server is chosen for application traffic delivery.

Summary

Multiple paths often exist between peers in most large enterprise and service provider networks. These paths can be used to reroute traffic in case of failures or link performance degradation, and can be utilized simultaneously to improve user experience through improved resiliency and higher throughput. Most traditional WAN optimizers and SD-WAN solutions do not follow a failsafe approach to deliver traffic. The 128T Session Smart™ Router combines optimized server heuristics, intelligent path monitoring, and lossless application delivery to form a failsafe delivery model that ensures application traffic is delivered despite failures. In addition, innovative quality of service and traffic engineering enhancements ensure optimal non-stop application performance and superior end-user experience.