128 TECHNOLOGY

RESILIENCY

CONTENTS

INTRODUCTION

The connected world has no tolerance for downtime. In today's digitized world, connectivity is the lifeblood of most businesses. A recent survey[1] found that 50% of employees in an enterprise lose access to important information during a downtime. About half of them resort to non-compliant methods of information sharing during this period. Downtime reduces employee productivity by 34% and results in nearly half a million dollars of revenue loss per hour to the enterprise.

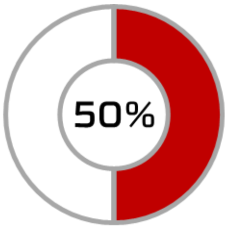

Employees
Lose Access
to Info

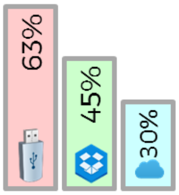

Employees
take Security
Risks

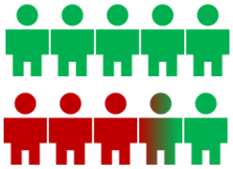

Employee
Productivity
Decreases

Revenue Lost
due to
Downtime

This also results in unquantified damage to reputation, customer loyalty, and compliance issues. 90% of organizations have lost access to critical systems and nearly a third deal with downtime every month.

The alarming frequency of these downtimes is due to the fact that legacy network gear tolerates and recovers from link and device failures without recovering sessions. They only provide high availability (HA) to minimize connectivity disruption through redundant components. Enterprises also independently pre-plan approaches to support services at an alternate facility for disaster recovery. These approaches do not provide zero downtime as they do not maintain any continuity of sessions.

The 128T resiliency solution provides virtually zero downtime by maintaining sessions though redundant clusters in a single or multi-site environment. It provides unprecedented elasticity through an N+M redundancy model, high reliability through fast failover by continuous flow state synchronization between appliances and innovative multi-site failover, and unlimited scale through hardware agnostic redundancy.

The solution builds upon 128T's Advanced Secure Networking vision which is rooted in five basic principles.

1. IP networks should be natively session-aware
2. Security and load balancing are not standalone functions
3. Routing must evolve to be application and service-centric
4. Deterministic routing, virtualization and segmentation does not require overlays
5. Zero-trust security must be omnipresent

[1] Devastating Downtime, Globalscape, 2014

The 128T resiliency solution ensures that the failover and switchover mechanisms are session-aware and completely secure.

128T RESILIENCY SOLUTION

The solution operates in Active/Active clustering mode. Multiple routers are grouped together as clusters, with multiple Active units processing traffic and sharing the network load. Each cluster node contains a minimum of two units acting as a Stateful HA pair. Active/Active clustering provides Stateful failover in addition to load sharing. The customer may choose to pass all traffic through one of the routers in the cluster. In this case the remaining routers in the cluster will not be processing traffic but they are all in Active mode with ability to process traffic if required.

The 128T solution operates in N+M redundancy mode where any number of routers participate in a cluster and they can act as backups of one or multiple routers in the cluster. Interfaces on different routers can be configured as redundancy groups. These redundancy groups are collection of resources that need to failover between the routers. An interface in a redundancy group is chosen as the primary and another as the secondary. This is done via a leader election or based on user defined priorities. Primary interfaces are used to route traffic through the cluster. In case of failure the traffic from the primary interface is switched to the secondary interface in the redundancy group via Gratuitous Address Resolution Protocol (GARP) or other routing protocol exchange.

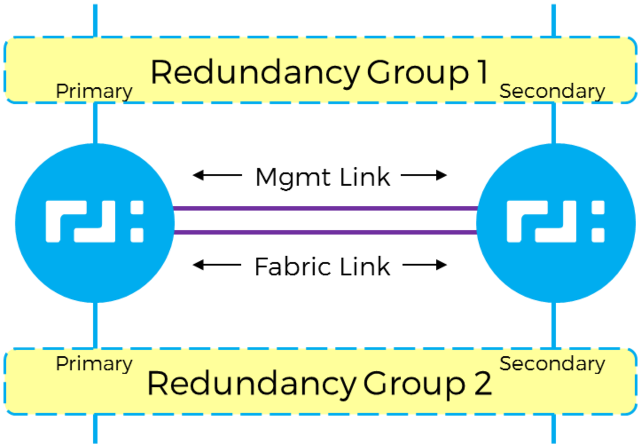

A fabric link between the routers is used to route traffic between them in case of failure. In the diagram these are shown as directly connected links but they do not have to be. Also the diagram shows two routers in a cluster for ease of understanding however there can be multiple routers in the cluster.

The management link between the routers is used to exchange routing and flow information between the routers. This is shown as a separate directly connected link between the routers in the diagram however it can share any link. All information between the routers are shared using highly efficient in-memory databases to minimize bandwidth usage and to enable instantaneous information exchange.

All processes in a 128T router are self-resilient. They can regenerate themselves independently in case of process failures or exceptions. Unless there is a dependency that requires other processes to restart or the device to switchover, the process will rebuild itself and establish communications to the existing processes. If a process failure requires another process to be restarted due to a dependency, then that process is restarted automatically. In-built self-checking mechanisms managed with software diagnostics ensure the integrity of the entire system.

The distributed nature of the 128T system and complete independence from the underlying hardware ensures that there is no limit of the number of routers that are part of a cluster. There is also no restriction on the number of interfaces that can be part of a redundancy group. This ensures that the solution is abundantly elastic. It can scale from a 1+1 configuration in a branch office to a fully distributed cluster in a large data center with N+M redundancy. This ensures a scale out architecture that can span numerous use cases and any possible scenarios.

SYNCING FLOW STATE

The 128T system syncs TCP and UDP flow state information between routers to ensure that no flows are lost and applications are not disrupted due to network outages. To ensure that no bogus or random packets cause the system to attempt to sync state, the 128T system only syncs state for established sessions. The 128T system also syncs all routing information to forward packets as traditional routers do in case of failure.

TCP Flow

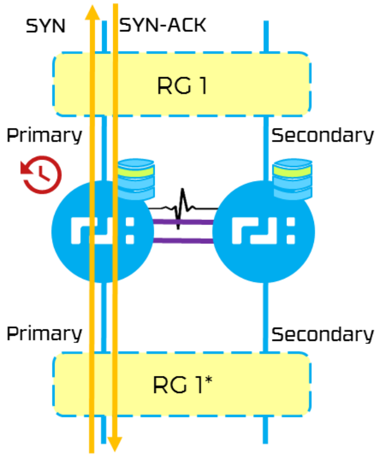

The 128T router creates an established flow record in its local in-memory database when a TCP flow is established. This is done when the 128T router receives a SYN-ACK packet in response to a SYN packet that it has forwarded recently. Once this record is created for that particular flow, it is synced with routers in the cluster so that interfaces in redundancy groups can take over traffic forwarding in case of failure. This state information is maintained as long as the flow is active. The record is removed when the TCP session ends or when an inactivity timer expires. All information related to the flow like policy, security, and quality of service rules are observed by all routers in the cluster to ensure complete reliability and security.

In the figure only two routers are depicted for ease of understanding. The 128T system can work with many routers in a cluster.

UDP Flow

The 128T router creates an established flow record in its local in-memory database when a UDP flow is established. This is done when the 128T router has forwarded a preset number of packets related to that UDP flow. Once this record is created for that particular flow, it is synced with routers in the cluster so that interfaces in redundancy groups can take over traffic forwarding in case of failure. This state information is maintained as long as the flow is active. The record is removed when the UDP session ends or when an inactivity timer expires. All information related to the flow like policy, security, and quality of service rules are observed by all routers in the cluster to ensure complete reliability and security.

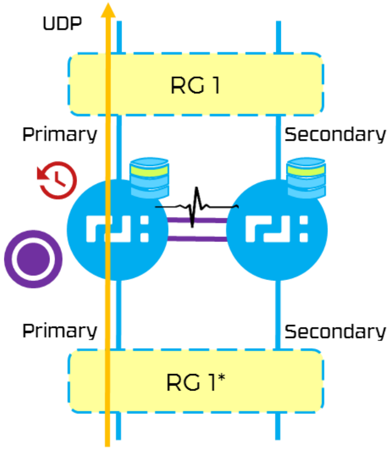

By intelligently identifying and syncing established flows, the 128T system is able to guarantee that end user applications do not see any interruption due to network failure.

MULTI-SITE FAILURES

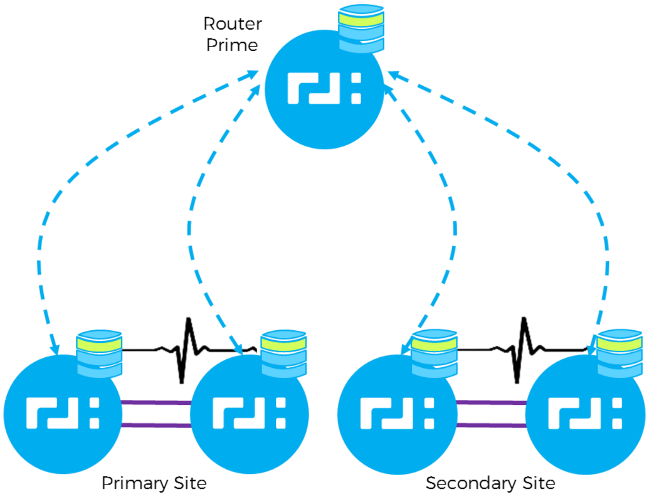

The distributed nature of the 128T system allows it to easily extend to multi-site failure scenarios. Router clusters in multiple sites can be configured as a collection. A router or router cluster configured as the Prime acts as a master store of flow information across the collection. In case of failure traffic can be rerouted via a secondary disaster recovery site. The secondary site can either query the master store on the Prime for information or the Prime can send the master store information to sites in the collection.

Another option is for the secondary site to send traffic directly to the Prime. The Prime already knows the flow information from its master store. It forwards traffic based on this information to the destination. This ensures that traffic flow is not interrupted while querying the master store and waiting for the information as is the case for many multi-site failover solutions today. If the flow information does not exist on the Prime, then the primary site had never seen this flow and it is just dropped.

On receiving packets back from the destination, the Prime forwards this to the secondary site from which it had received the packets. Once the secondary site gets these packets from the flow, it creates an appropriate record for this flow. The Prime can also respond back to the secondary site with the flow information if it does not receive any return packets. The secondary site can now forward packets for this flow directly to the destination without forwarding it to the Prime.

FAILURE SCENARIOS

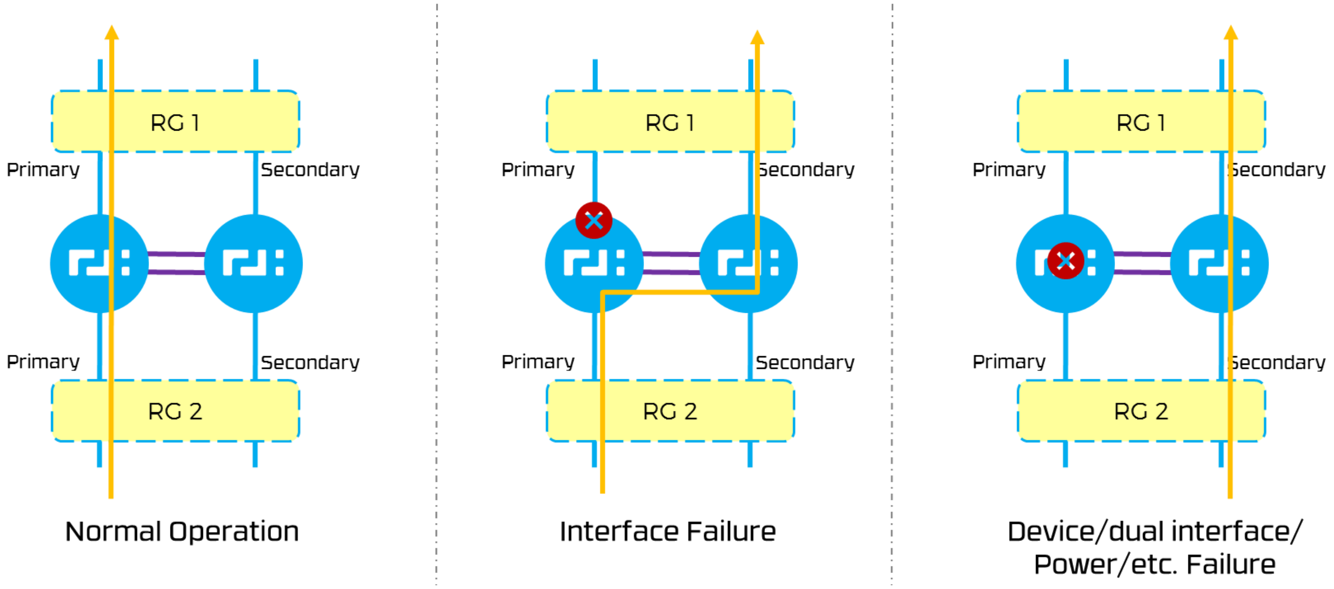

The 128T resiliency solution can enable zero downtime failure protection for all types of planned and unplanned network outages. It ensures that the end user application is oblivious to any network outage. The solution can provide high availability protection and switchover in case of process, interface, component, device, and cluster failovers. These may be caused by power outages, human errors, or other factors.

In case of normal operation all traffic is forwarded through the primary interfaces of the redundancy groups. These may be on a single router or may be spread across the router resulting in different flows to go through either router.

In case of a single primary interface failure, traffic will be routed through the fabric link to the secondary interface on the other router. This fabric link does not have to be a dedicated directly connected link. In case of process, device, component, or dual interface failures that completely disables the router passing traffic then the traffic switches over completely to the other router.

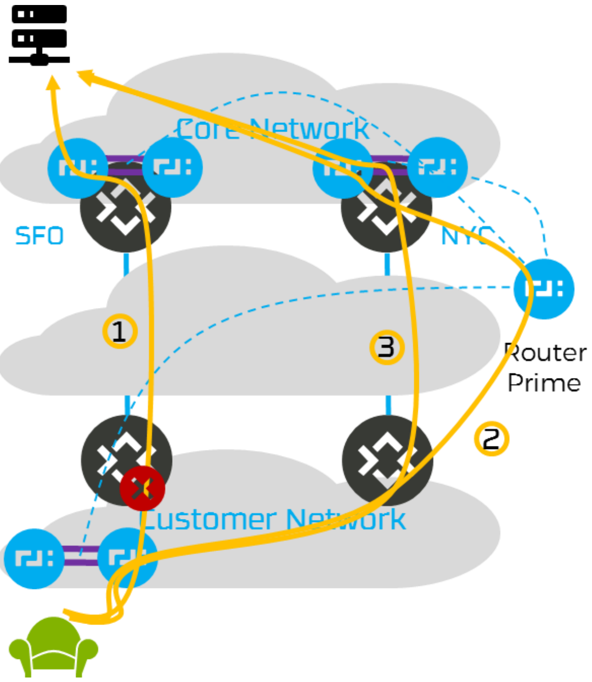

The solution can also handle disaster recovery across multiple sites in case connectivity to an entire site or the complete path is lost. Flow 1 shows the path the flow takes during normal operation. In case of failure which renders the path impossible to use, the flow is sent to the Prime which forwards Flow 2 to the correct destination through an alternate site. After a response is received or the Prime informs the source of the updated path, the flow depicted as Flow 3 is directly sent to the alternate site. This ensures that no single path or site failure will disrupt the end application if paths through multiple sites exist ensuring disaster recovery.

IN SERVICE SOFTWARE UPGRADES

The 128T resiliency solution enables the ability to support in-service software upgrades. This reduces downtime due to planned upgrades. The 128T system provides information to switch traffic to different routers in a cluster. Any router within the cluster can be isolated to prevent traffic forwarding through it while other routers in the cluster continue to forward all the traffic. This isolated router can then be upgraded. The upgraded router can then continue to participate in the cluster. Traffic can be switched back to this router or it can remain dormant in the cluster while fully Active to process traffic if required.

SUMMARY

Legacy routing vendor solutions only provide 1+1 redundancy with high availability for connectivity. They are unable to maintain flow information so any change in path or NAT will break the end-to-end application connectivity. They do not offer perfect multi-site failover. They require overlays, additional protocols, expensive equipment, complex networking, and additional licensing costs to offer disaster recovery options across sites.

The 128T resiliency solution offers the industry’s most comprehensive, advanced, and innovative high availability and disaster recovery solution along with stateful failover. This guarantees virtually zero downtime, multi-site failover, and a scale out architecture with N+M redundancy. This is a must-have for organizations to maintain connectivity during planned and unplanned downtime. It prevents revenue loss, improves productivity, reduces security risks, improves customer satisfaction, and guarantees regulatory compliance.

The 128T system ensures complete peace of mind by keeping businesses running at top speeds all the time.

128 TECHNOLOGY

www.128technology.com | info@128technology.com