Users of strategic applications are often assured that their network will be perfect. As users continue to discover, even perfect networks can still fail critical needs.

5 Ways That Your Perfect Network Fails

Contents

Introduction

Local area networks (LANs) are considered perfect for most applications. They transport data quickly and reliably. Moving from a LAN to a wide area network (WAN) for strategic applications can be more challenging. WANs are often complex networks that don't provide ideal data transport. Still, some WANs are extremely robust and are considered perfect. What is a perfect network?

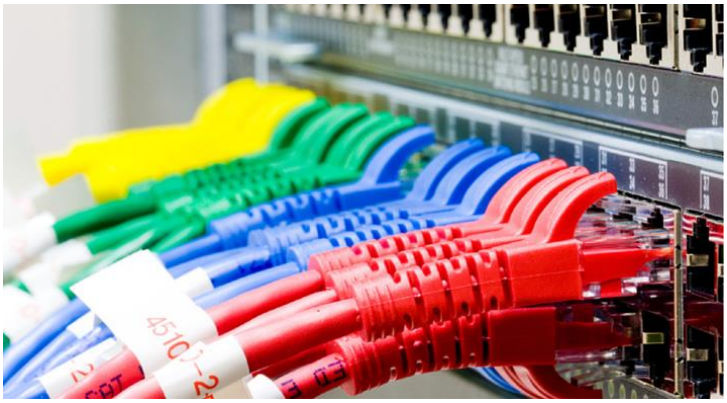

4 Attributes of a Perfect Network

When you are assured that your network is perfect, it should:

- ✓ Never discard or corrupt packets
- ✓ Never delay, reorder or duplicate packets
- ✓ Never limit throughput
- ✓ Impose no relevant limitation on packet size or rate

Even perfect networks aren't perfect all of the time. There are situations when a perfect network fails your mission. This happens when the perfect network stops being perfect, possibly for a short but high impact time, and often in an unexpected way.

How do perfect networks fail strategic applications? This eBook identifies 5 ways that you may not have thought about.

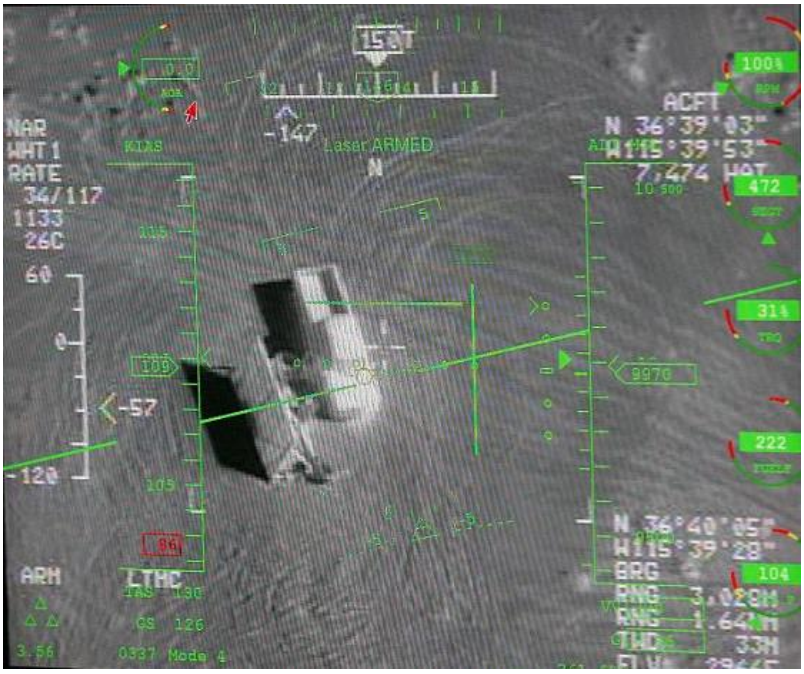

1. High Round Trip Time (RTT)

The time it takes a packet to move from one end of the network to the other and back again is called the Round Trip Time (RTT). Due to the laws of physics, WANs that traverse long distances have relatively high RTTs, and this impacts performance by interacting with the network protocol in two ways. The most commonly used transport layer on IP networks is Transmission Control protocol or TCP. When TCP encounters a high RTT, it deliberately transmits data much more slowly than the network potential. This is called Slow Start and occurs each time an application begins transmitting data after a pause. The result is that in the face of high RTT, TCP adds even more delay. The following chart illustrates the effect of Slow Start on a high RTT WAN.

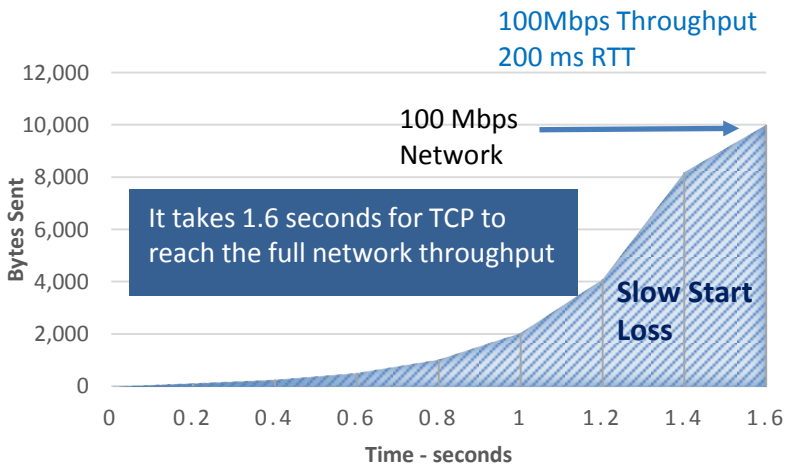

TCP Doesn't Like Long Fat Networks

Slow Start

TCP uses slow-start to avoid sending more data than the network is capable of handling. It does this by sending a small number of packets to begin and gradually increasing the amount it sends as it receives responses. If it fails to receive a response, TCP determines that the maximum rate has been reached and slow start is over. Therefore the slow start algorithm is sensitive to the length of the round trip time.

Bandwidth Delay Product

For a network link to be fully utilized, the amount of data on the link at any given time must equal the link bandwidth multiplied by the link delay. This value is known as the Bandwidth Delay Product (BDP).

TCP's Receiver Window (RWND) is a fixed parameter that controls the quantity of data that can be sent before the sender stops to wait for an acknowledgement. When RWND is less than the BDP, the transmitter cannot fill the line and available bandwidth goes unused. Thus, for a given RWND and link bandwidth, a longer link delay means less data can be sent.

TCP also responds to high RTT networks by limiting its maximum transmission rate. When the Bandwidth Delay Product exceeds TCP's Receive Window, TCP cannot saturate the link, and valuable bandwidth is wasted. TCP interacts poorly with high RTT networks, presenting challenges for high bandwidth or latency sensitive applications. These are explained more completely in the sidebar on the previous page.

Long Fat Networks

Networks that span multiple continents have long round trip times regardless of network quality or potential bandwidth. This is due to the transit time for electrical or radio signals and can not be avoided. A rule of thumb is that a 5,000 mile network will have a round trip time of at least 100 ms just due to distance, regardless of router behavior.

Network	WAN RTT	Maximum TCP Rate Mbps
Denver to Washington, DC	45 ms	11.7
New York to Hamburg	86 ms	6.1
London to Sydney	294 ms	1.8

Transferring Satellite Data Across WANs

Due to its poor response to high RTT (latency), TCP limits data rates for transfer across satellite networks. The inherently long RTT limits throughput. The figure below shows how TCP constrains data rate due to distance for various satellite types.

Standard TCP windows are set to 64KB. While there are extensions that permit larger windows, they are often not supported by operating systems or firewalls. On long latency networks, this window size drastically restricts the usable bandwidth. For instance, if a T1 transmission line of 1.5 Mbit/second was used to transport data from a satellite link with a 520 millisecond round trip time (RTT), the link will only run at just over 1 Mbps or 68% of the possible maximum.

For this reason, strategic applications designed for satellite or other networks with high latencies use UDP instead of TCP. However, UDP provides no repair mechanism or guarantee of delivery. Many networks are designed for file transfer, VoIP and other more forgiving applications. For UDP users, is the network still perfect? Strategic applications are faced with a dilemma: use TCP and tolerate low bandwidth and long delays, or use UDP and suffer data loss.

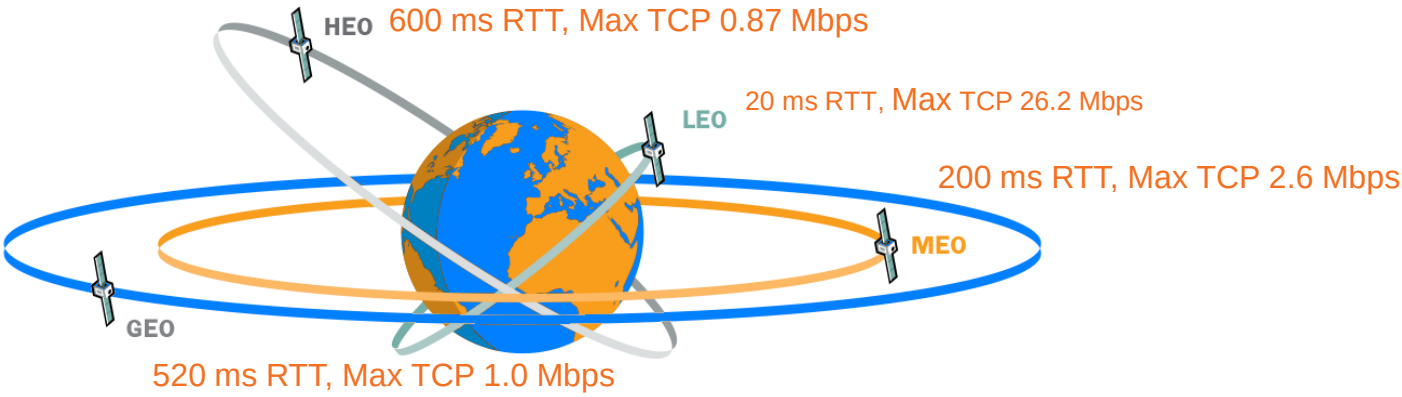

2. Misconfigured Equipment

A perfect network doesn't lose packets or corrupt them, but network performance and reliability can be affected by misconfigured equipment. Yankee Group estimates that 62% of IP network downtime is due to configuration errors. "Network managers don't have a good way of checking configurations, and the impact on network security and reliability is not immediately felt."

For example, configuring a router with an incorrect Maximum Transmission Unit (MTU) will cause packet fragmentation. This increases the packet rate and bandwidth of IP streams, which may result in loss. Additionally, fragmented streams are more likely to suffer reordering.

Because configuration changes are a routine part of network maintenance, a previously perfect network link is always at risk of being degraded. The result is that strategic applications relying on perfect packet delivery may fail without warning.

Understanding that frequent configuration changes will be made, and that these changes can lead to packet loss or the possibility that insufficient bandwidth will be allocated, there is always a possibility that these problems will occur in the future. In short, even if you don't experience these impacts today, it's reasonable to worry about it.

Encrypted Network

Even perfectly configured equipment can periodically drop data. Encrypted networks use secret keys that may only be used for a limited amount of time to protect a limited amount of data. Once these keys expire, the secure tunnel must be rekeyed, and this process often causes network dropouts. Dropouts will lead to holes in the data stream that may be unacceptable for strategic applications.

3. Equipment Failure

Perfect networks fail suddenly when something goes wrong. Examples from recent history include a cable that was cut by vandals, a truck that crashed into a bridge eliminating a core link, and fan failure resulting in a router shutdown. In all of these cases the perfect network abruptly ceased being perfect.

When a network loses some portion of its capacity, the immediate response is to move traffic onto the remaining capacity. This means that even when your link isn't degraded by the failure, you can lose provisioned bandwidth to rerouted traffic. Your perfect network is no longer perfect for the duration of the event.

It doesn't take an unforeseen disaster to take down a link. Network transceivers can fail, causing bit errors to creep in. Single bit errors cause checksums to fail, resulting in the loss of the entire frame.

IT Superhero

Losing a WAN link creates a high priority support issue. Disaster recovery calls for all hands on deck. Traffic will be rerouted as quickly as possible. However, even a brief outage can be disastrous for strategic applications, and since these events happen more often than expected, a measure of insurance is prudent.

Police told the newspaper that four fiber-optic cables were severed shortly before 1:30 a.m. PDT along Monterey Highway. A telecom spokesperson confirmed that the fiber that was cut appeared to be the work of vandals. But she explained that fiber cuts happen all the time, typically due to an accident. "Fiber cuts happen more often than people think," she said. "Usually it happens accidentally when someone is drilling in the ground, landscaping a lawn or repairing some other infrastructure in the ground. We know this happens all the time, so we're ready to reroute traffic whenever we have to."

What Does an SLA Really Mean?

Mission critical networks specify important parameters like uptime, latency, and packet loss. Even if the SLA performance is acceptable for your application, what does it mean when a failure occurs? The network operator refunds your fees for the time when the network is out of compliance. Seriously? You need your data on time, not your budget refunded. In many cases the SLA levels are actually much lower than the typical performance. SLAs are supposed to be guarantees after all. What if the network operator actually delivers down to their SLA – can your application handle it?

4. Congestion

Example of Abrupt Packet Loss on a Normally Reliable Network

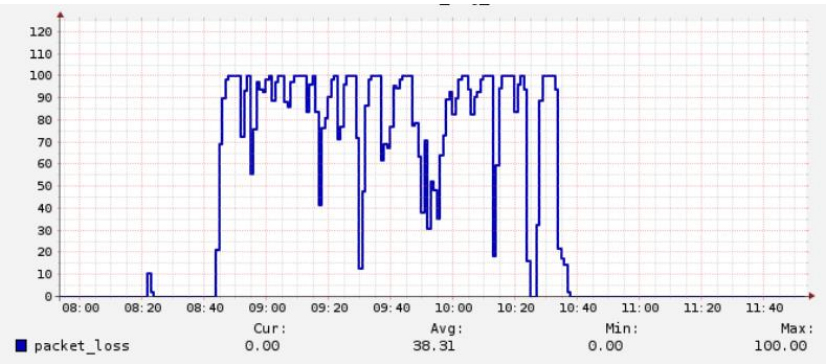

When a network is perfect, people want to use it. Invariably, the demand and the traffic increases over time. When network traffic exceeds capacity, core routers drop packets, resulting in data loss.

An increase in users isn't the only cause of increasing congestion. Applications can misbehave and exceed their bandwidth allocations. When they do, the extra traffic can bleed over into your bandwidth. Applications and users naturally exceed their allocation gradually over time. This happens as more users discover more resources to be accessed. From time to time, specific events can cause network traffic peaks.

An additional factor is understanding performance dips that occur. If previously perfect networks are no longer satisfying application needs, it can be difficult to know why without sufficient analytic tools and IT support..

Data Growth

According to Cisco, annual IP traffic will reach 3.3 ZB by 2021. In 2016, IP traffic was 1.2 ZB. In case you haven't run across a zettabyte yet, a zettabyte is 1000 Exabytes. An Exabyte is 1 billion Gigabytes (GB).
These numbers are too big to really understand, and admittedly they are global traffic, not your organization.

Still, the expectation is for data volume to almost triple over the next 4 years.

The basic drivers of this traffic will affect all organizations. Network congestion is growing and hardware investment is unlikely to keep up. Even if performance is terrific today, a measure of caution is in order and it may be useful to think about insurance against this data tsunami.

5. Redundancy

One of the things that makes a network reliable is redundancy. When a link fails, traffic is rerouted to another path. However, this path switch causes a brief interruption in data flow. This is inconsequential for file transfer or web browsing, but is catastrophic for strategic applications that require data to be delivered in a continuous and timely way.

Further, the law of unintended consequences may play a role. Network redundancy is based on the idea of eliminating a single point of failure, usually by having active-active links. In this case two paths are continuously active, providing load balancing features and fail over to one another in the event of failure. After failover, the link that carries 100% of the traffic is likely to behave worse than its normal performance. Strategic applications will see a drop in performance due to the unexpected congestion.

.

Who Could Be Against Failover?

The priority for network redundancy is to keep the data flowing in the event of failure, This can certainly be beneficial for applications that have their connectivity severed. The cost however, is born by applications that maintained their connectivity, In the event of network failover, your performance may change dramatically.

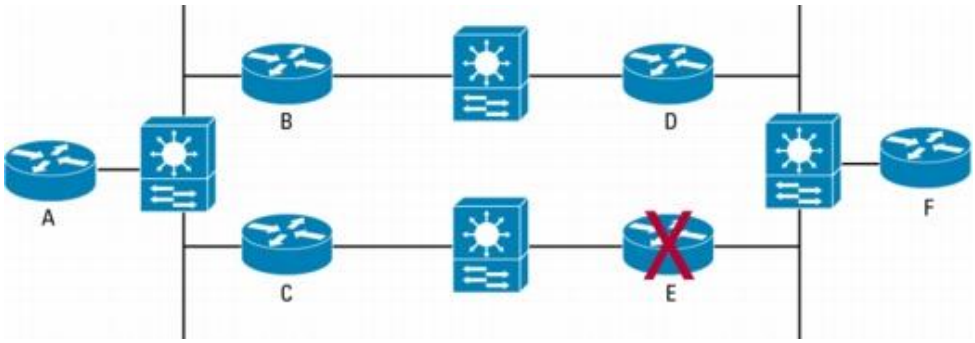

There are more extreme cases. When one link fails, part or all of the service goes down. The reason is that despite having two paths, complex configurations sometimes have dependencies that are unknown and untested. Even more extreme cases involve a measure of bad luck. While doing maintenance on one link, its redundant pair fails. Another version is believing there is redundancy but finding it wasn't configured correctly. The most embarrassing case is when failure results in a resource shortage, resulting in cascading failures. This happens from time to time, resulting in large scale national internet outages that make the news.

When done correctly, redundancy reduces the probability of a total failure but does not assure unaffected performance. When done incorrectly, redundant solutions may be less reliable due to increased complexity.

This eBook outlined five ways that a reportedly perfect network can fail you. In an increasingly connected world, network performance issues are a growing problem. To learn more about the root causes of these issues, please see the next piece in this series, *10 Reasons your WAN is Broken, available at www.rtlogic.com/products/datadefender.*

DataDefender guarantees performance of strategic applications that run across complex networks

Phone: 719.598.2801 • Email: sales@rtlogic.com • Web: www.rtlogic.com