



WHITE PAPER

Evolving Networks leased lines

Redefining failover

evolving.net.uk

Evolving Networks leased lines

Redefining failover

1 Introduction

There is a widespread assumption that leased lines don't fail, and that, should the unthinkable actually happen, solid SLAs and "guarantees" mean they will always be up and running again without damaging delay.

These are myths. Everything fails. Sooner or later, today, tomorrow or in three years' time, whatever the technology may be, it will fail. That's as true of leased lines as of anything else. And, as we will see, no SLA can guarantee uptime or restoration time in the event of a fault.

This white paper explores the issue of how you can make your leased lines truly resilient, ready for any failure.

Sooner or later,
today, tomorrow
or in three years'
time, whatever
the technology
may be, it will fail

2 Leased lines have an ugly secret

The assumption that leased lines are somehow exempt from the universal principle that all technology will fail is down to the widespread misconception that should a fault arise on a leased line, its SLA guarantees its restoration within a fixed time period.

This simply isn't the case. Far from guaranteeing restoration of a faulty line within a certain period of time, leased line SLAs merely specify what will happen — usually a penalty payment to the customer — should the line not be restored within a given time period.

That's a critically important distinction. Which network manager wants to tell their MD that the business has been offline all morning, and when it will be back online is unknown, but there's no cause for alarm, because the leased line provider will credit them a paltry amount in compensation?

SLAs in fact provide no real guarantees whatsoever in terms of uptime or fault fix times, and their prescribed penalty fees will rarely come anywhere near the costs to the business of even brief outages. For organisations dependent on their leased lines, this is serious indeed.

SLAs in fact
provide no real
guarantees
whatsoever in
terms of uptime
or fault fix times

3 Resilience means failover?

Talk about resilience and it's likely that the word "failover" will be thrown around a fair amount. The concept of switching over from a failed resource to a good one is absolutely central to most resilience strategies: it's almost always assumed that resilience means failover.



This assumption is erroneous and limiting. In this white paper, we explore a better way, specifically with reference to leased lines. It's an approach that doesn't leave resources hanging around, unused, in reserve for that rainy day when primary resources fail, instead allowing full use of all resources at all times, relegating none to the role of backup.

At the same time, it can deliver levels of resilience limited only by the customer's ambitions and budget, as well as additional performance, traffic control and network visibility benefits.

It can deliver levels of resilience limited only by the customer's ambitions and budget

4 The solution is technological, not contractual

SLAs don't resolve the problem of leased line failure because no amount of contractual verbiage can change the simple fact that technology fails.

Perhaps counter-intuitively, though, better deployment of technology can mitigate its own tendency to failure, to as great or little a degree as is desired in any given situation. Intelligent software has a significant part to play in this.

The principle is familiar: identify and remove single points of failure, wherever they exist.

A leased line can never provide resilient connectivity on its own, because, as a single circuit, it is a single point of failure. In fact, it is several single points of failure, lined up one after the other, a problem in any one of which can take it offline. These boil down to two main areas: the cable connecting the customer to their provider's core network infrastructure, and that core network infrastructure itself.

Such single points of failure must be removed, by adding redundancy — a further circuit or circuits, delivered by multiple providers, with core network infrastructures completely separate to each other.

A leased line can never provide resilient connectivity on its own, because, as a single circuit, it is a single point of failure

5 "Resilience" isn't always resilient

Many providers offer resilience options for leased lines, but, typically failing to address all single points of failure, they are often of limited benefit.

For example, many "big name" providers offer two circuits for resilience, but deliver their service through a single core network infrastructure. The additional resilience provided by a second cable will be of no help at all in the event of a failure in the shared core network.

At the customer site, those two cables will often be brought into the building through a single wall duct. Damage to that part of the building could take both cables out — another single point of failure.

Furthermore, those two circuits are in almost every case just that: two separate circuits. The idea is to use one of them, keeping the second in reserve in case the first fails.

Typically, the first inkling a customer will have about such a failure will be the site going offline. They must then disconnect the failed line, connect the reserve line and reconfigure IP addresses, firewalls and other systems and services.

Even where failover is more automated, it will rely on a measurable service interruption. This will impact end users, for example dropping important sessions, business transactions and voice calls. It may still involve a change of IP address.

However failover is achieved, there will be no guarantee that the backup circuit will work when it is brought online. It must, of course, be paid for, though, month-in, month-

Many "big name" providers offer two circuits for resilience, but deliver their service through a single core network infrastructure



out. A costly resource, it may never actually be used, and may not deliver at all in the customer's hour of need.

Such issues are not being addressed by traditional ISPs – for the solution, we have to look above the ISP.

6 Better than failover: transcend the network

Evolving Networks have been delivering the solution to this problem since 2008: the multipath aggregation of connectivity.

Our Intelligent Network Fabric (INF) combines the bandwidth of two or more circuits (of any type) presenting them as a single connection with the aggregated bandwidth of the multiple physical circuits that make it up.

Operating on a software and management layer above any single ISP, the INF always uses all the bandwidth available in all constituent circuits in each connection, so there is no investment wasted. When a fault arises in a circuit, restricting or even entirely removing its bandwidth, the INF simply continues to use all of the available bandwidth – that of the remaining sound circuit or circuits. No failover is needed, and no interruption is experienced by end users.

To address the core network infrastructure single point of failure, Evolving Networks operates the UK's first Software Defined Network, spanning multiple geographically diverse datacentres in a critically decentralised way unheard of in the industry. Circuits within any given connection connect to completely separate core network infrastructures and network operators.

The result is a connection which vectors packets around faults in individual circuits or core network infrastructures without actually needing to route around them at all. In most cases the customer is completely unaware that a fault has occurred until advised by Evolving Networks.

The INF makes almost any level of resilience achievable, for the first time allowing the customer to decide for themselves how resilient they need their connections to be.

7 From myth to reality: resilient leased lines

Business should not fall for the smooth sales patter of providers who pretend their solutions are resilient when they actually hide various single points of failure, nor the lie that SLAs guarantee anything beyond the (typically wildly inadequate) compensation payments to be made in the event of their breach.

The hard truth is that leased are not inherently resilient, any more than any other single circuit is, and their SLAs cannot guarantee uptime, performance or fault fix times.

By managing network traffic at the packet level, through software and above any single ISP, exceptional resilience, uptime and performance is delivered. This is achieved not by promises in legalese but through the solid, proven engineering excellence on which every Evolving Networks connection is built.

Contact Evolving Networks today to find out more about resilient leased line connections and how they can give your business the uptime and bandwidth you need.

Evolving Networks have been delivering the solution to this problem since 2008

CASE STUDY Arthur David



Arthur David needed resilient leased line connectivity at their three sites in Bristol and Somerset.

Read how Evolving Networks delivered the perfect solution in the Case Study.



Evolving Networks
Nexus House
7 Commerce Road
Lynch Wood
Peterborough
PE2 6LR

+44 330 55 55 333

sales@evolving.net.uk

evolving.net.uk