Network Availability of Competing Wave Service Architectures

Architecture Choices for Enterprise Network Architects

September 30, 2019

Will Russ, Product Manager, Wavelength Services

Version 1.3

Table of Contents

Purpose

This document provides guidance for achieving the highest levels of network availability on a domestic DWDM transport backbone network. The guidance describes state-of-the-art trade-offs between costs and performance that drive critical architecture decisions when selecting protection schemes. The network availability of competing architectures is quantified to predict their performance in deployment when compared in simplified, relative *classes* of overall cost. These cost classes can impact architecture selection wholly, when costs must be contained. The unique and industry-leading analysis of availability performance vs. cost also guides the customer on which types of additional redundancies (diversity vs. protection) are most effective when an increase in network availability is desired to support increased resiliency.

Scope

This document details the design and configuration of the customer's transport network to include point to point Wavelength Services over the Verizon network, with or without third party leases.

An overview of available protection schemes and the trade-offs between them is given.

A detailed analysis of the network availability of the Data Center Interconnect (DCI) for each architecture option is provided. It includes the availability of the carrier's equipment and customer's equipment to calculate the end to end availability of the entire Data Center Interconnect, including all redundancies. This measures the availability from the customer's perspective to include the customer's routers.

Detailed configuration of other equipment is beyond the scope as well as network management architecture or maintenance parameters.

Verizon Wavelength Services overview

Verizon's Wavelength Services provides high speed dedicated point-to-point and point circuits between Customer Sites on the Verizon Network. Services are provided and are available in Metro, National and International configurations, subject to availability on our network and through our partners. Wavelength Services are available from 1 Gb/s to 100 Gb/s speeds today. The standards-based customer interfaces include Ethernet, OTN, SONET, SDH and SAN protocols at various speeds. Support for route diversity, equipment diversity and site diversity provide resiliency while protected wave service options increase the availability per circuit. Edgeless access options prevent carrier technician access into customer sensitive buildings. Layer 1 encryption solutions provide an added layer of security. Latency SLAs provide assurance of your application performance.

Quantifying the resilience of highly available transport & routing architectures

Protection switching schemes and router redundancies are the key architecture attributes

Resiliency can be increased by adding redundancy to various points in the transport network and performing failover switching at layers 1, 2 or 3. The deployment of protection switching or network re-routing at these different points is the key aspect that separates competing architectures (that have similar costs) that each aim to increase the overall network availability between two data centers.

What makes certain architectures "competing"?

Most IT Network Architects don't deploy a resilient transport architecture using unlimited funds. Comparing an architecture that is twice as expensive as another isn't fair, since you can increase redundancy with expenditure. It's therefore more useful to compare architectures with similar costs, whose ranges fall within the same orders of magnitude. Two architectures whose costs are similar are therefore competing and each has different levels of network performance for achieving resilience. That performance is quantified mathematically using predicted failure rates and calculating the network availability as the metric.

Network availability quantifies the resiliency of an architecture

Measuring resiliency for a transport network circuit (in this case, a Data Center Interconnect) is well understood to fall within the Reliability Engineering discipline. Using the standard approach, the quantification of an architecture's resiliency is done by statistically predicting the reliability of the DCI circuit's *network availability*. Network availability is expressed as the ratio of uptime per unit of time (as a statistically predictable average of performance over time).

Predicting the network availability using a mathematical model is a reasonable exercise, since the (statistically) predicted failure rate of equipment can be obtained directly from equipment vendors and fiber outage rates can be estimated by empirical data.

Network availability calculations MUST include the routers

It's not enough to simply maximize availability of the carrier's circuits. The overall availability of the customer's point to point Data Center Interconnection (DCI) that includes the routers is the ultimate metric, since the routers are part of the Wide Area Network (WAN). In this study, the availability is calculated to include the customer's router at the data center, not just the carrier's equipment. For the Routers, the Router Chassis, fabric and WAN side port modules are included. The router ports facing the Data Center's switching equipment and/or servers is excluded, as it's the same for all architectures. For the Wave services transport, the Modules, fabrics, chassis and fibers are included as components:

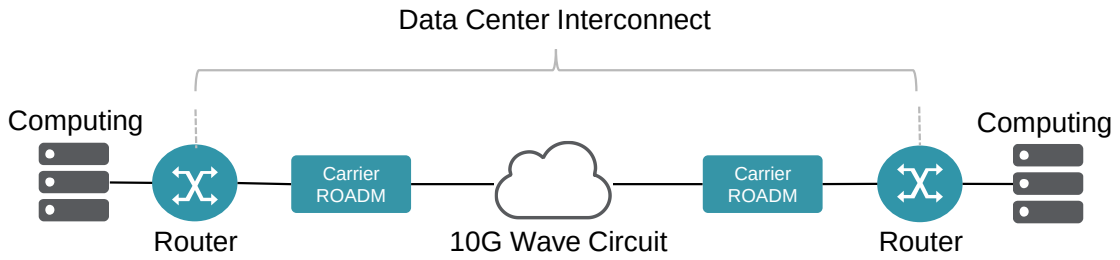

Figure 1: The entire WAN data communications path is calculated, not just the Carrier's wave circuits

Four classes of competing architectures

When moving beyond a single link as Data Center Interconnect in order to increase resilience, the competing architectures fall within four ranges of relative cost (monthly recurring charges and/or non-recurring charges) to the customer. The Classes that compete (have similar equipment costs and wave service costs) are:

- Dual "Waves" from a Single Router (Single Protected ckt vs. Dual Unprotected Circuits)
- Four "Waves" from a Single Router (Dual Protected Waves vs. Four Unprotected Waves)
- Dual Routers (Dual Prot Waves vs. Four Unprotected Waves, with Dual Routers)
- Transport Network Element Equipment Diversity (Dual Prot Wave vs. Four Unprotected Waves with Dual Routers, both of which include Network Element Equipment Diversity)

Modeling the hypothetical reference circuit for the "Effective Data Center Interconnect"

The actual hypothetical reference circuit is not given but is described to the reader as "typical" for a 1600 km long haul 10G wave circuit with customer routers from the enterprise class. While different vendors and different designs can alter the results slightly, the relative performance differences between architectures remains intact for design variations. The designs are objective and realistic, as the carrier's performance is not the objective. Since the reference circuit contains both carrier and customer equipment, the relative performance of the entire *effective* data center interconnection among the architectures is the objective. The effective data center interconnect is the ability to transmit at least 10G of Ethernet data between data centers. A pair of 10GbE links deployed that experiences the failure of one of them is still an effective DCI under this study. If the Enterprise architect sees 20G of bandwidth demand between the data centers as the minimum requirement, then the bandwidth is simply doubled but the relative performance of each architecture remains the same.

Even though the modeling details cannot be shared, they are made with the intention of withstanding the most scrupulous of reviews to remain objective and realistic. In most cases, the explanation of the results for each architecture's performance will become intuitive to the reader.

Key modeling assumptions the network designs are "real world"

It's of critical importance to note that the assumption for the diversity between the architectures with four waves (either 2 protected circuits x 2 channels per circuit or four unprotected circuits) is a key factor in the outcome of these higher end architectures. In this study, we assume the four waves are routed over only two geographically diverse routes instead of four. That models a network design that chooses lower latency over resiliency, if a summary is made.

Another modeling assumption is that OTN switching equipment is used, which grooms 10G circuits into 100G trunks. The route diversity assumed in these architectures does not include redundant OTN switch nodes until the "Network Element Equipment Diversity" option is invoked.

This guidance is without liability to Verizon

This guidance is simply a statistical prediction of performance and cannot be warranted and does not guarantee service quality or component reliability. It is voluntarily offered as general strategic guidance to our customers without consideration. It must be considered an estimate and not actual and therefore must be taken without liability to Verizon. The customer is encouraged to seek alternative sources of reliability engineering information.

The source of failure rate data

The FIT rates for the routers and ROADM (Re-Configurable Optical Add-Drop Multiplexer) modules and chassis are supplied by the vendors (state of the art metro and ULH ROADM industry volume leaders are chosen, as is a typical industry leading Enterprise Router). The fiber is assigned a MTBF using historical fiber outages. Estimated fiber outages from Proactive Maintenance of inline amplifiers and ROADMs is also included. Processor modules, power supply modules and switching fabrics are assumed to be redundant on all equipment and those with traffic affecting status are included in the model (shelf processors are not). The availability of the DC Power feed is assumed to be 100% (via battery backup and generator backup), so it is not modeled.

The role of cost in the choice of architecture

The Network Architect can provide circuit redundancy, diversity and failover switching (layer 1 protection switching or layer 3 rapid re-convergence) with various combinations of both as an architecture. Each architecture drives a certain number of router ports and a certain number of "waves". A "wave" (defined loosely here for this network availability study only) is either a wave circuit proper (say a 10GbE) or it could be a protect channel or working channel in a single protected wave circuit. They both require two wavelengths through the carrier's network, so they're similar in cost (from a Wavelength Services perspective). It's wise to understand which expenditures (on the equipment and wavelength services for an architecture) achieve more availability for your buck. Each of these competing architectures is illustrated below.

Note that one protection scheme (IXC Mesh Restoration) can reduce the cost for wavelength services of the protected option significantly, when compared to its equivalent architecture on the unprotected side within that class. To understand why the mesh restoration reduces the cost of the wavelength services for the equivalent number of waves in the unprotected side, recall that mesh restoration will deploy an ample of amount of spare capacity on the OTN mesh network but that spare capacity may be shared by other parallel working paths and therefore other customers of the carrier. So while the amount of spare capacity deployed always protects against any single fiber outage, it can be shared among other users, which reduces the cost per user. This savings is significant in a highly meshed network. This becomes significant because this savings, coupled with lower router port costs for the protected side's architecture make the protected side less expensive than the unprotected side in each of the classes. So while the unprotected side performs the best (because layer 3 is protected additional router ports) it also costs more. The question is, how much better does it perform? Are you getting bang for your buck by adding more router ports, or is the performance improvement of layer 1 protection sufficient?

Modeling assumptions

- The network availability calculated is the "entire data center interconnection" which includes the Enterprise router chassis and WAN side ports and the carrier's wave service transmission equipment. This approach is required to measure the impact of varying router port counts in various architectures
- The availability of modules and chassis (shelves) are calculated from FIT rates (Failures in Time) given by the equipment vendors. Leading ROADM and Router vendors are used and configured in a typical long haul circuit configuration of 1600 km using metro ROADMs, ULH (Ultra Long Haul) DWDM transport and the OTN Mesh
- 4 hour Mean Time to Repair (MTTR) is assumed
- Fiber diversity is assumed for the unprotected pairs and for the working and protect channels of the protected circuits. The diversity of the "four unprotected" circuit architectures is assumed to be across four diverse paths in the IXC segment but are limited to two diverse paths in the access sections, to model a worst case scenario for metro market connectivity
- Empirical fiber cuts per year are used and a value for outages during proactive maintenance is used for the unprotected circuits. This includes force majeure events
- The IXC portion of the hypothetical reference circuit is assumed to be protected by mesh restoration where the spare capacity deployed will support all single fiber outages and some secondary fiber outages. The modeling reflects continued restoration of the protect channel, so that it essentially never fails from fiber outages (protect channel fiber availability is assumed to be 100%)
- Switching time during protection is ignored, as it's statistically insignificant for long term averages for downtime

- The FIT rates and availability of the modules, chassis and fiber outages aren't given, as they're considered extremely sensitive competitive information. Only the results of the calculations are given.
- The availability of each component (modules, chassis, and fiber outages) is accurate to 30 decimal places.
- The classes are identified as simply those architectures which consume an identical number of "waves" (either an unprotected circuit or a working or protect channel within a protected circuit) and share similar router and transport equipment costs. Costs are not given.

The analysis of the results for each class is given:

Class 1: Two "waves" from a single router

Two Unprotected circuits (a diverse mated pair) competes with a single protected circuit as far as cost of the wavelength services goes, so these two compete head to head. The single protected circuit is a single Router link and a single wave service but recall that its working and protect channels are two "waves" across the network, making its cost similar (but less than) that of the two unprotected circuits (which also require two waves):

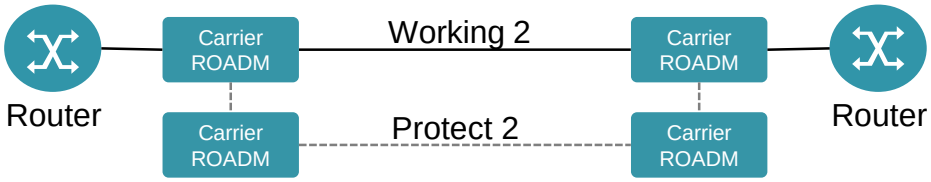

Figure 2: One protected wave circuit. Single router ports but with Fiber & ULH equipment protection

The single protected circuit's design lacks redundant router ports and that penalizes its performance. Its downtime is a hefty 190 minutes per year. The failure rates of the router ports is significant and that shows here, when they're left unprotected. The layer 1 protection used is SNC protection in the metro and IXC mesh restoration for the IXC segment of the wave service.

On the other hand, the [competing] unprotected pair of circuits gains redundancy for the router ports but doesn't have fiber protection against cuts or proactive maintenance events or protection of the ULH transport equipment:

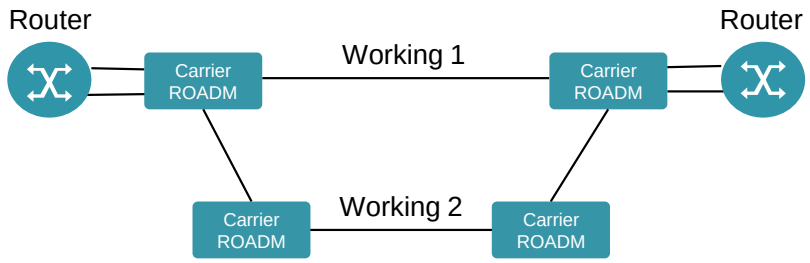

Figure 3: Two unprotected wave circuits. Protects router ports but not Fiber & ULH equipment

While the competing pair of unprotected circuits provide much better performance, remember it comes at the cost of additional router ports and possibly a router chassis upgrade. Nonetheless, this pair of diversely routed unprotected links is the benchmark for many enterprises today.

Class 2: Four waves from a single router

This class improves beyond the pair of unprotected circuits (2 waves) with either:

- a pair of protected circuits (4 waves via 2 ckts) or
- by using four unprotected circuits (also 4 waves)

For the protected option (2 protected circuits), this adds additional resilience over class 1 by now achieving redundant router ports (because its two circuits instead of one). The layer 1 protection used is SNC protection in the metro and IXC mesh restoration for the IXC segment of the wave service.

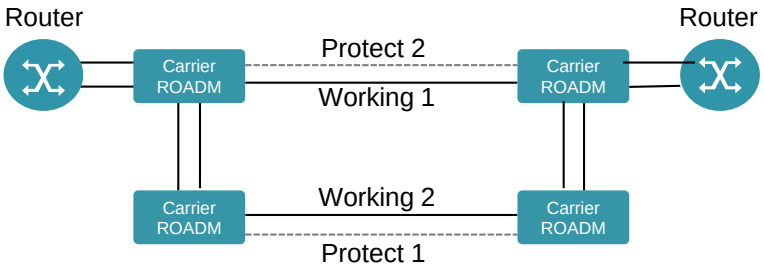

Figure 4: Dual protected wave circuits. Adds redundant router ports and Fiber & ULH equipment protection

Note that the performance is nearly identical to its competitor (albeit by coincidence, in the face of so many variables). Compared to its competitor (four unprotected waves), having equivalent performance makes the "protected pair" option a bargain:

- since its router costs are lower (two ports instead of four at each router) and
- since its wave service costs are also lower (recall that two protected waves are cheaper than four unprotected diverse waves when mesh restoration is used).

As for its competitor, the four unprotected waves, it quadruples router port redundancy (as well as metro ROADM transponder redundancy) and even quadruples fiber protection in the IXC via four unique routes but the cost is high, for both the wave services and for the routers (via ports and maybe even fabrics). The recommendation at this level of spending is therefore to use dual protected circuits above. An illustration of the four unprotected circuits:

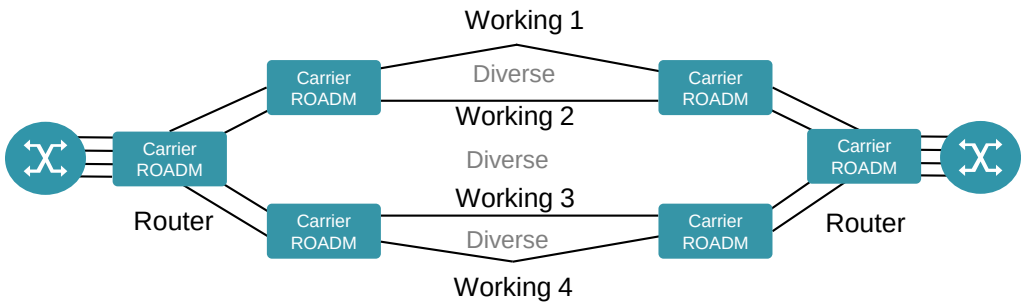

Figure 5: Four unprotected wave circuits via a single router. Quadruples router port, Fiber & ULH redundancy

Class 3: Dual routers

This class eliminates the router chassis (and some common equipment) as a single point of failure which improves the performance significantly. The two circuits still use a single ROADM chassis at the Customer premises and again at the LD POP and the OTN switch chassis is also still a single point of failure. This is considered a minimum level of resiliency for enterprises with active-active computing structure across multiple data centers, where the DCI is transactional.

The protection option is for two protected wave circuits from the dual routers:

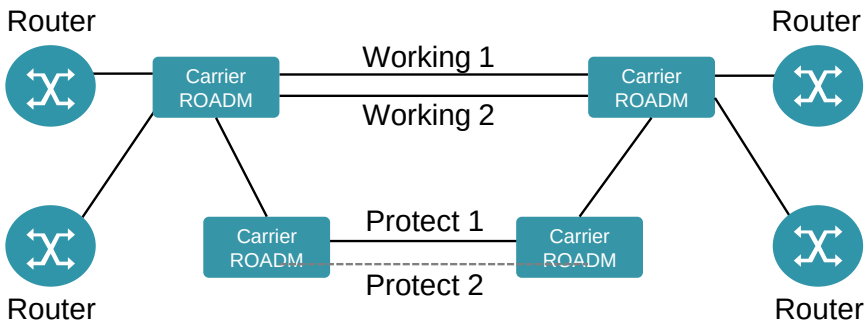

Figure 6: Two protected wave circuits via dual routers. NE equipment diversity eliminates ROADM chassis and OTN switch chassis as single point of failure

The mesh restoration, used as protection for the IXC portion of the circuit, not only provides protection against a working fiber outage but would also restore a second outage and even a third, as long as there's spare fiber capacity deployed. Again, the performance is almost identical to its competitor, so again, the recommendation is to use the protected wave architecture for two protected waves vs. four unprotected.

The competing architecture (based on class of spending) is four unprotected wave circuits from the dual routers. While this quadruples router port redundancy and fiber and ULH redundancy, the quadrupling of those failure rates isn't a bid improvement over simply two redundant components:

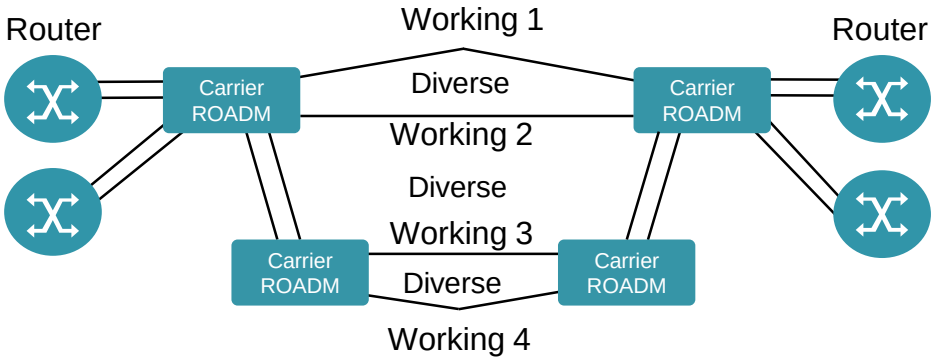

Figure 7: Four unprotected wave circuits on dual routers. Quadruples router port redundancy

Class 4: (Transport) Network element equipment diversity

The ultimate in network architecture, this eliminates the Metro ROADM (and/or NID) chassis at the customer premises and at the LD POP as a single point of failure as well as the OTN Switch chassis as a single point of failure. The protection option uses two circuits. The two protected circuits' working channel is routed into separate Network Elements (DWDM Transmission systems or optical NIDs) for the entire length of their path. Note that the protect channel of one circuit may still ride on the same equipment as the working channel of the other circuit:

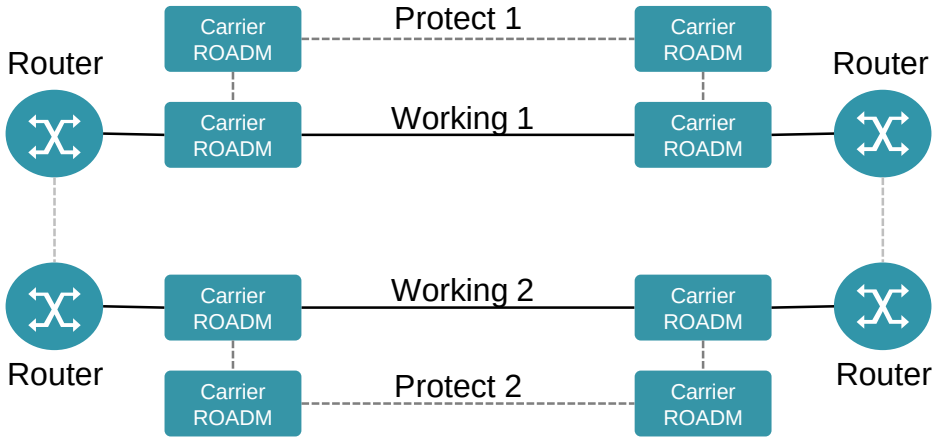

Two Protected Wave Circuits via Dual Routers

NE Equip Diversity Eliminates ROADM Chassis and OTN Switch Chassis as Single Point of Failure

Figure 8: Two protected wave circuits via dual routers. NE equipment diversity eliminates ROADM chassis and OTN switch chassis as single point of failure

The performance is essentially the same for this vs. the competing architecture, which is four unprotected circuits over the NE equipment diversity, so since the protected option is less expensive for wave service and for router port costs, it's the recommendation for the ultimate in resiliency.

The competing architecture is shown (four unprotected circuits over diverse network elements):

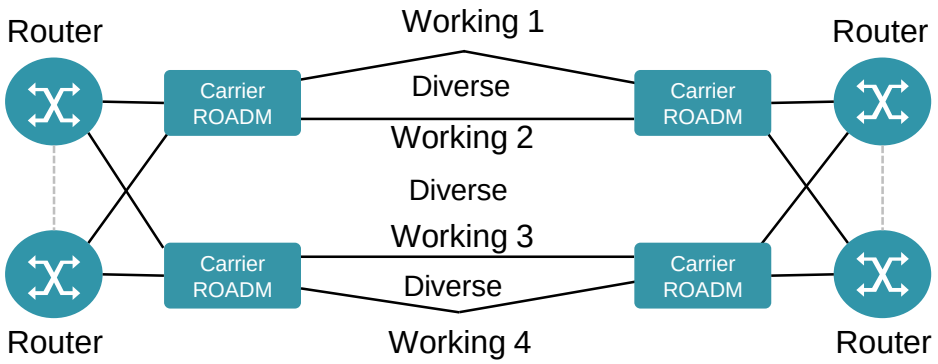

Four Un-Protected Circuits over Diverse Network Elements

Figure 9: Four unprotected circuits over diverse network elements

The resulting network availability for all four classes

Each progressive architecture eliminates one or more single points of failure (a module, chassis or fiber) to reduce the downtime. While it's easy to simply select the architecture with the best performance, remember that this comes at a cost, so many architects will need to consider the more moderate options and get bang for their buck.

Each of the protected wave options are in the same "cost" ballpark as their unprotected competitor within a class. Cost here is defined as the carrier's cost for fiber and equipment to provide the architecture. The assumption is that cost will be proportionate to price, so the theory is that these classes are representative of wavelength services prices to the customer in orders of magnitude.

The network availability of each architecture is graphed below. The network availability here is the percentage of uptime where at least one link in the architecture remains up, so it is in effect a statistical average of the effective uptime:

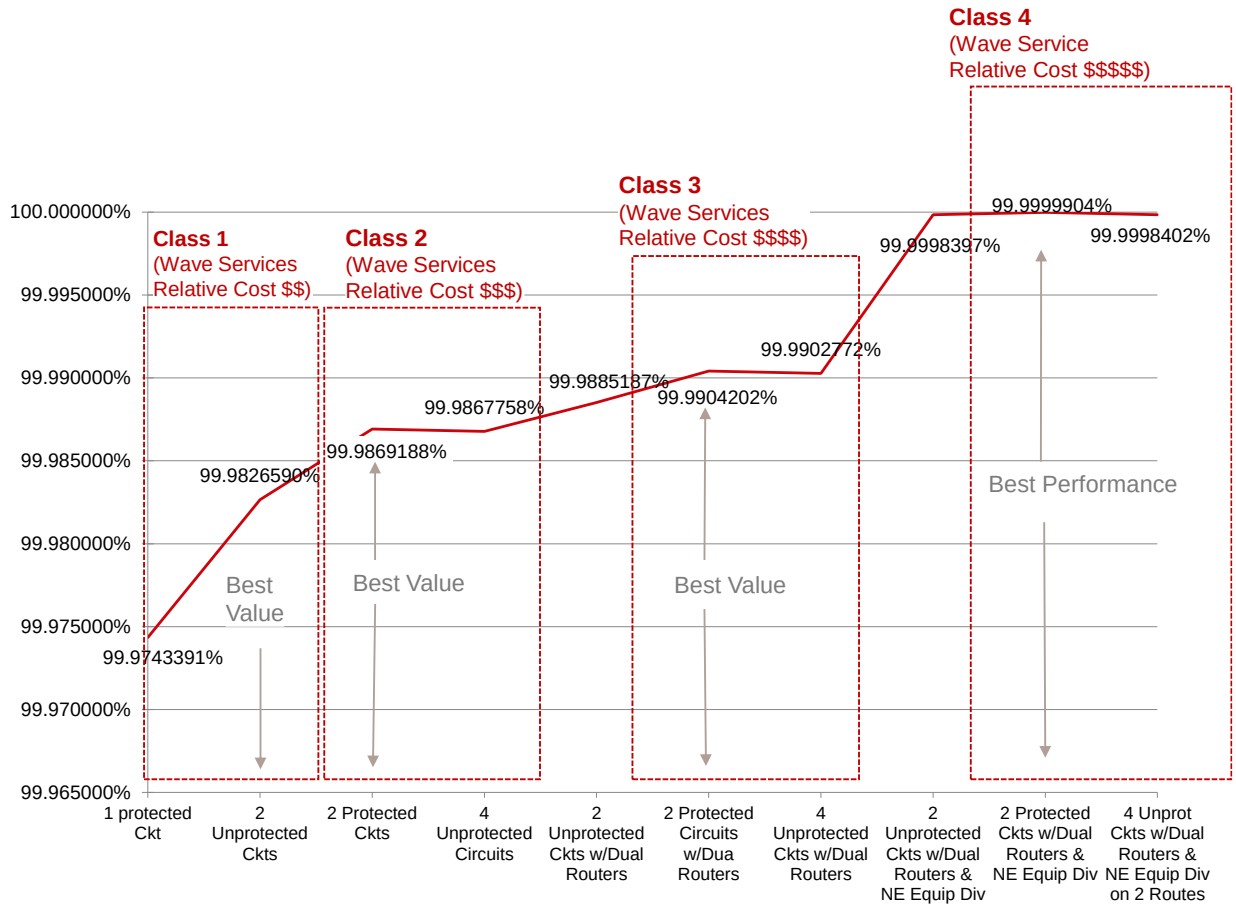

Figure 10: Network availability (statistical average) where at least one circuit is up. Includes all wave fiber and equipment (including chassis) and router (chassis and ports) failure rates

The graph below shows the same results measured as the reciprocal, which is considered more intuitive (shows the respective downtimes for each network availability percentage. Recall that this is a statistical average per annum for each architecture:

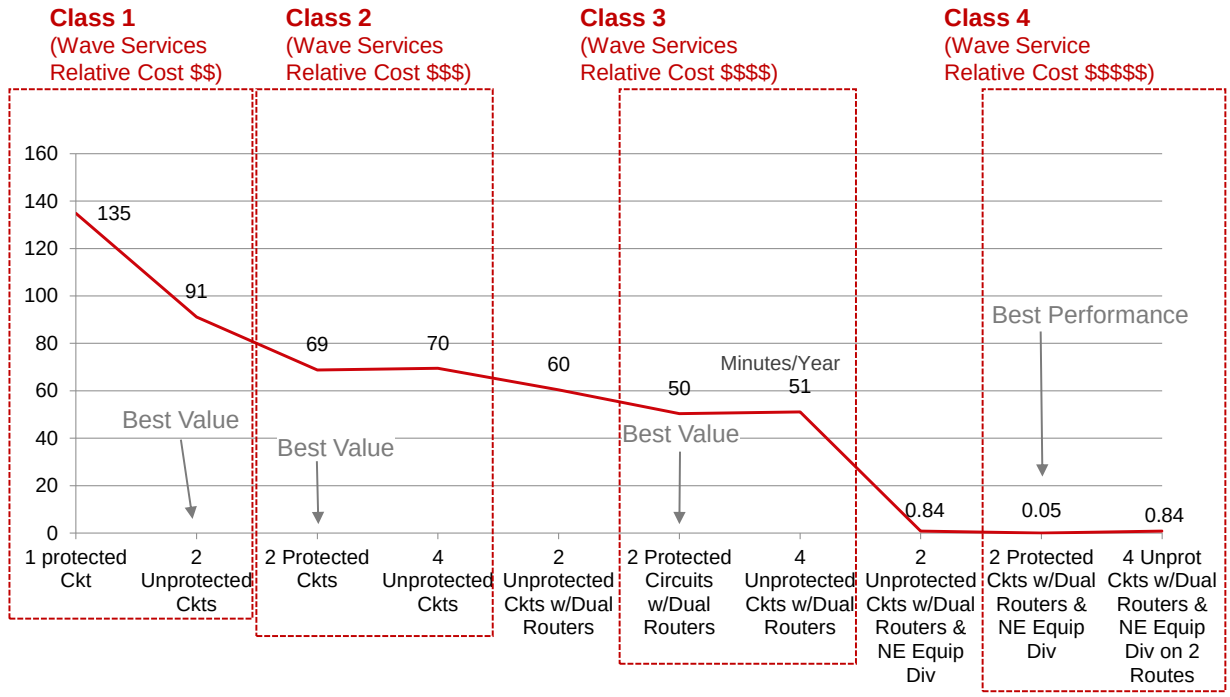

Figure 11: Downtime per year in minutes (statistical average) where all parallel circuits are failed. Includes all wave fiber and equipment (including chassis) and router (chassis and ports) failure rates

Conclusion

For class 1, it's always best to go beyond a single protected circuit and use two unprotected wave because it adds redundancy for the router ports. This "diverse mated pair" can be thought of a table stakes.

Improving the availability performance beyond that of the diverse mated pair in class 1 is best done by deploying two protected waves instead of four unprotected waves. This is because the wave services cost is identical and the router costs are far lower (two ports on each end rather than four). The performance is also slightly better.

Improving beyond class 2 is best done by deploying dual routers at each end. This creates redundancy in the router chassis, power supplies and other common equipment. The chassis is the key component there.

Improving beyond class 2 is best done by adding Network Element Diversity to the Wave service. This creates redundant ROADM chassis throughout the carrier's network and eliminates them as single points of failure. This is the absolute best performing architecture.

The recommended highly available architecture

Using the above analysis, the architecture using dual Protected Wave circuits on dual routers and Network Element Equipment Diversity is the recommended *highly available* architecture:

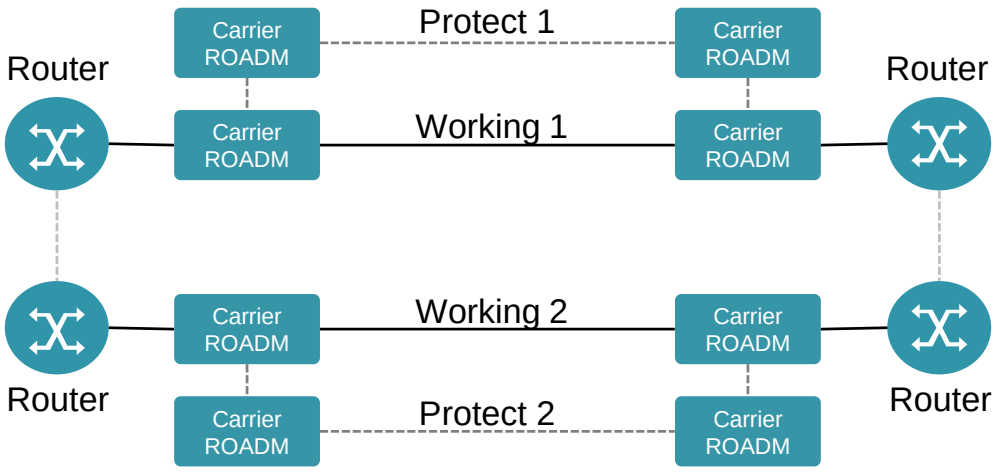

Figure 12: Two protected wave circuits via dual routers. NE equipment diversity eliminates ROADM chassis and OTN switch chassis as single point of failure

This architecture includes the Wavelength Protection feature combined with the Network Element Equipment Diversity feature. Note that the monthly cost of the protected option is lowered when the carrier provides the protection via Mesh Restoration, which shares the spare capacity among multiple parallel working circuits to reduce the cost of the protection bandwidth. This is a key differentiator among carriers.