WHITE PAPER

infinera®

CORIANT IS NOW PART OF INFINERA

COST-OPTIMIZED STRATEGIES FOR MAXIMIZING TRANSPORT NETWORK AVAILABILITY

Delivering the Highest Availability at the Lowest Cost for Each Service

As more and more services become mission-critical, as enterprises demand ever more stringent Service Level Agreements (SLAs), and as consumers become less tolerant of disruptions to their services, high network availability has never been higher on the agenda. Examples include the 5G PPP targeting "zero perceived" downtime and a major North American carrier stating its new target as "six nines" (99.9999%) availability.

In an ideal world where cost is no object, every service would have close to 100% availability. However, in reality higher availability normally comes at a cost, and not all services the transport network needs to support have the same requirements. Some services are more cost sensitive and more tolerant of downtime. Other services will rely on other layers to deliver the levels of availability required. One size does not fit all. This white paper explores strategies to address the challenge of providing the highest level of network availability at the optimal cost for each service.

THE DEMAND FOR HIGH NETWORK AVAILABILITY

The need for high network availability has never been greater. Driven by applications including mission-critical e-health and process optimization in manufacturing, the 5G PPP is targeting 99.999% availability and “zero perceived” downtime. In addition, one major North American operator, at the 2016 Big Communications Event, announced that it was aiming to improve its offering by guaranteeing customers “six nines” availability.

As residential broadband becomes commonplace not just for browsing the internet but also for online gaming, communications, and video, disruptions become more painful and increase the likelihood of end-users switching to another provider. Likewise, a stringent SLA including availability, jitter, packet delivery, and latency was ranked the number one factor influencing Ethernet service purchasing decisions by enterprise and wholesale customers in Heavy Reading’s September 2015 Carrier Ethernet Survey.

AVAILABILITY	DOWNTIME PER YEAR
99%	87 hours, 40 minutes
99.9%	8 hours, 46 minutes
99.99%	53 minutes
99.999%	5 minutes, 16 seconds
99.9999%	32 seconds

ONE SIZE NETWORK AVAILABILITY DOES NOT FIT ALL SERVICES

However, higher network availability has a cost, and not all services require the same levels of high availability. Some services are less mission-critical than others. For example, remote meter reading does not require the same level of availability as backbone energy grid communications. Remote surgery needs much higher availability than the routine backup of medical records. Cost sensitivity is also a factor: the premium for high availability might be trivial to an investment bank but a significant additional cost to a Small/Medium Enterprise (SME).

The other factor driving the levels of availability the transport network must deliver is which layers are used to deliver the required level of availability. Services that rely on only the IP layer or a combination of transport and IP layers to increase availability will have less stringent requirements from the transport network than services that rely purely on the transport network to provide high availability. Likewise, an enterprise or wholesale customer may supply their own high availability by taking unprotected transport services from more than one service provider and providing their own resiliency mechanisms over the top.

THE POTENTIAL CAUSES OF TRANSPORT NETWORK DOWNTIME

Key events that can trigger network downtime include fiber cuts, equipment failures, site failures, human error, and fiber degradation. Long repair times also have the potential to exacerbate network downtime.

FIBER CUTS

A key cause of network downtime in many geographies is fiber cuts. Fiber cuts can occur accidentally during construction, road work, or maintenance of a utility such as electricity, gas, or water. Natural disasters including earthquakes and hurricanes can also lead to fiber cuts. Additional causes include deliberate sabotage or vandalism. Vandalism might be mindless, such as taking potshots at overhead fiber with a rifle, or involve an intentional economic motive, such as incorrectly believing the cable contains valuable copper and attempting to steal it.

The frequency and causes of fiber cuts will vary from geography to geography. In developed markets, one cut per 100 km per year might be expected, while in a dense metro region in a developing market experiencing rapid economic growth, multiple fiber cuts on a daily basis are not uncommon.

TRANSPORT NODE FAILURES

The transport nodes themselves can fail if critical common functions such as power, fans, fabrics, or controllers fail. Likewise if interfaces or port cards fail, any services they are carrying could be impacted. Each module of the transport node will have a mean time between failures. For active hardware, this will be measured in tens of years to hundreds of years depending on the complexity and level of integration. For passive components such as filters, this will be measured in thousands of years and can largely be ignored.

SITE FAILURES

Even if there are no failures in the transport equipment, a failure in the data center or central office housing the equipment, such as loss of power or a failure of the cooling system causing the network equipment to overheat, could occur and could cause the node to fail.

HUMAN ERROR

Installation and configuration errors can also lead to network downtime. Commonplace stories abound of how a careless elbow inadvertently took down traffic to a major customer or between two major cities. Poor maintenance, including the failure to change fan filters on a regular basis or to maintain the required operating temperature in a central office, can also be a source of network downtime.

FIBER DEGRADATION

As discussed in the Coriant white paper *Maximizing 100G+ Reach in Long Haul Networks with Challenging Fiber Conditions*, factors including splices and fiber additions made during repair, bending, dirty connectors, and environmental conditions including temperature, humidity, and pressure can change the properties of fiber over time, thereby increasing attenuation and dispersion. For longer and more challenging wavelengths that were engineered with a limited amount of margin, eventually this margin could be diminished resulting in unreliable performance or downtime.

LONG REPAIR TIMES

Downtime is a function of both frequency of failures and time to locate faults and complete repairs. The speed of the repair is a function of how long it takes to be aware of the failure, how long it takes to identify the location and cause of the failure, how long it takes to get the relevant people and equipment on-site, and how long it takes to make and test the repair. The quantity and location of spares and test equipment can strongly influence repair times for certain types of failures.

OPTIONS FOR MAXIMIZING NETWORK AVAILABILITY

NETWORK-LEVEL RESILIENCY MECHANISMS: PROTECTION

Network-level resiliency mechanisms can protect against a wide range of failures including fiber cuts, the failure of other nodes or sites, and interface and port card failures, so they are often the key focus for delivering high availability to the network infrastructure as well as to individual services.

Transport protection typically involves only data plane mechanisms for fault detection, and recovery is typically very fast with a gold standard of 50 milliseconds recovery after the failure has been detected and the need for action has been determined. Furthermore, all the decisions about the protection path are calculated and planned well in advance of any failure. Protection can be implemented at different layers and against different types of failures. The most common type of transport protection is 1+1 with traffic sent on both the working path and the protection path, with the receiver choosing the best quality signal, enabling very fast fail-over. Alternatively, 1:1 protection has a working path and a protection path on which traffic is not sent until after the failure has occurred. 1:1 protection requires signaling between the two ends of the circuit to coordinate the fail-over, so it is inherently slower than those 1+1 schemes that do not require signaling (i.e., unidirectional), though it can still be very fast. M:N protection provides a variant of 1:1 protection that uses M number of protection paths to protect N number of working paths, a good example of which is active-standby LAG (Link Aggregation), which could have N active links and M standby links.

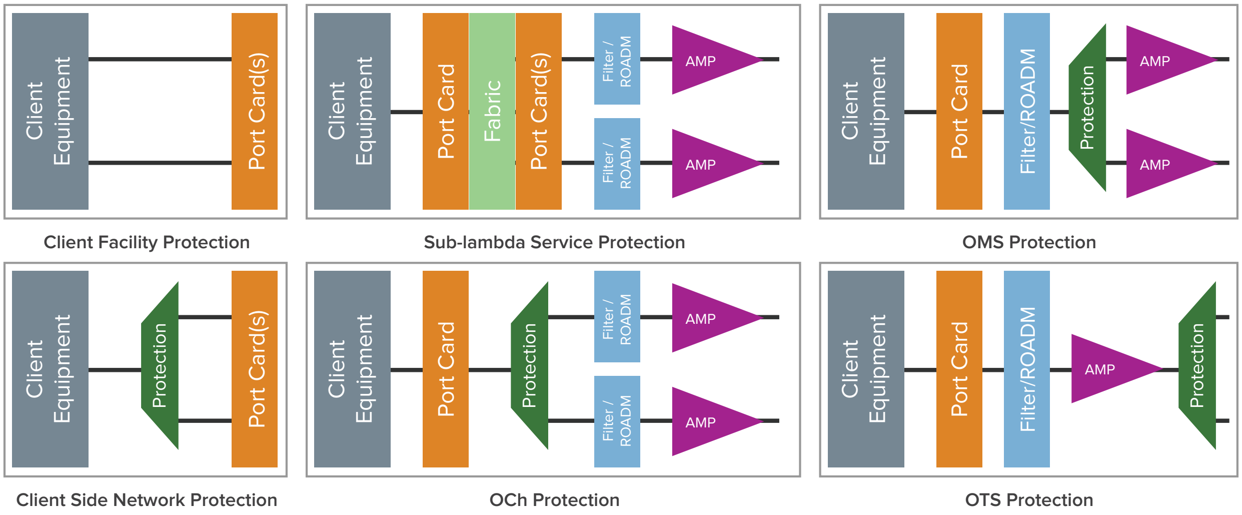

Client Facility Protection

Sub-lambda Service Protection

OMS Protection

Client Side Network Protection

OCh Protection

OTS Protection

FIGURE 1 – Key Network-level Protection Options

Common protection types in transport networks include:

- **Client Facility Protection:** Protects against a failure of the interconnection facility between the client equipment and the network equipment, including the client equipment ports, network equipment ports, and interconnect cables. Examples include 1+1 APS for SONET, 1+1 MSP for SDH, and Link Aggregation (LAG) for Ethernet. It provides the highest level of protection when combined with network protection or restoration but comes at a higher cost.

- **Client Side Network Protection:** Protects against a failure on the network equipment the client port connects to, with the client equipment port connecting to two ports or modules on the network equipment, together with network protection of the end-to-end path interconnecting the client ports. Examples include passive Y-cables and active protection switching modules. Being purely passive, Y-cables benefit from a mean time between failures measured in thousands of years, while active protection switching has the advantage of being able to support protection with port cards in different systems, as the protection decision is taken on the active protection switching module, and therefore, there is no need for communication between port cards. Client side network protection has a higher cost than network-only protection, but unlike client facility protection, it does not protect against a failure of the client equipment port itself.

- **Sub-lambda Service Protection:** Protects against failures on line interfaces and across the network for individual client services (ports or interfaces) or high level sub-lambda containers. Common mechanisms include SNC for SONET, SDH, and OTN, and G.8031 VLAN protection for Carrier Ethernet. Sub-lambda service protection enables different levels of protection for different services and typically delivers very fast protection, however, unless combined with either of the client protection schemes listed in the first two bullets, it does not protect against client side failures.

- **OCh Protection:** Protects against the failure of an individual wavelength by taking the DWDM interface, duplicating the signal, typically with a splitter, and sending it down two paths and then using an optical selector at the other end. OCh protection can provide cost-effective protection of end-to-end wavelengths without expensive duplication of DWDM interfaces. However, the DWDM line interfaces are now single points of failure, and it can be challenging to deliver very fast recovery with coherent colorless (see the section, "The Challenge of Coherent Colorless Protection").

- **OMS Protection:** Can be used to cost effectively protect all the lambdas on a fiber against a fiber cut or failure of the amplifiers on an Optical Multiplex Section between two DWDM add/drop nodes. However, it does not protect against other network equipment or client equipment failures or provide end-to-end service protection. Obtaining very fast recovery for coherent channels can also be challenging.

- **OTS Protection**: Can be used to provide very cost-effective protection of all the lambdas on a fiber against a fiber cut between two adjacent nodes. However, it only protects against fiber cuts on a single span and requires similar fiber lengths on both spans.

THE CHALLENGE OF COHERENT COLORLESS PROTECTION

While 50 ms protection switching is standard for 10G based on direct detect, coherent transmission common with 100G+ creates some additional challenges for fast protection. First, if the protection path has different properties regarding attenuation, chromatic dispersion, PMD, etc., it can take the DSP some time to tune to this new set of parameters.

The second challenge relates to detecting failures when using coherent colorless add/drop. Technologies used for colorless and colorless/directionless add/drop include splitter/combiner and star couplers, while optical multicast technology is used for CDC (Colorless/Directionless/Contentionless). As all of these schemes rely on the receiver seeing multiple channels and then tuning to the correct one, loss of signal can no longer be used for fast failure detection.

NETWORK-LEVEL RESILIENCY MECHANISMS: RESTORATION

Restoration differs from protection in taking action to establish the backup path only after the failure has occurred. The path itself can be calculated after the failure has occurred (dynamic restoration) or before the failure has occurred (pre-planned restoration). The principle benefits of restoration are the ability to survive multiple failures, the ability to more efficiently share protection resources amongst multiple services, and the ability to take a more optimal protection path using minimal additional resources. However, it is important to be aware that the ability to survive multiple fiber cuts will be constrained by the network topology and the degree of “meshiness” in the network.

Traditional ASON/GMPLS restoration schemes have relied on mature IP/MPLS protocols including OSPF-TE as the routing protocol and RSVP-TE for signaling, with extensions for transport. Disadvantages include slower recovery times relative to data plane based protection, the risk that the required resource may not be available, and additional complexity. Recovery times for ASON/GMPLS restoration will depend on the layer, with Layer 1 restoration typically much faster than Layer 0 (optical) for which interfaces may need to retune, ROADMs may need to open new paths, and amplifiers may need to rebalance. However, it is possible to combine protection and restoration with protection providing guaranteed 50 ms fail-over for the first failure with the control plane establishing an alternative backup path so that the next failure can again be met with 50 ms recovery. This approach incurs additional cost but can provide the 50 ms recovery time of protection with the ability of restoration to survive multiple failures.

However, as the industry moves toward SDN as an architecture for management and control, SDN can provide an alternative to ASON/GMPLS enabling a more open and centralized solution for restoration.

Key benefits of SDN include multi-layer and multi-domain restoration and overcoming the challenge of interoperability between vendors, which remains a challenge for ASON/GMPLS.

As the distributed nature of ASON/GMPLS still has its advantages relative to the more centralized approach of SDN, architectures are emerging that combine ASON/GMPLS for distributed fast restoration and SDN for multi-layer, multi-domain restoration or for services that do not require the fast restoration of distributed ASON/GMPLS.

HIGH NODE AVAILABILITY

High node availability can be achieved by eliminating single points of failure in the common equipment, including power feeds, power supplies, fans, controllers (i.e., system processors), and fabrics. While this provides a cost-effective insurance policy against failures, it cannot protect against failures to port cards, fiber cuts, site failures, or human error, so it is typically only part of the solution rather than a complete solution in itself.

As a more cost-effective alternative to redundant 1+1 common equipment, functions such as the fabric and power can be distributed to individual modules. Where centralized fabrics are required for electrical switching scalability 1:N (1 redundant, N working), fabric protection can provide a more cost-effective alternative to 1+1, as the percentage of cost allocated to protection is reduced as N increases.

PROACTIVE DOWNTIME PREVENTION

Protection, restoration, and high node availability all rely on detecting and responding to failures after they have occurred. However, many faults such as increased attenuation, reduced OSNR, increased bit error rates, higher equipment temperatures, or harder working fans may not cause a failure straightaway. A complementary approach that can deliver even higher levels of availability is to identify potential problems and take action to resolve them before they can impact availability.

This strategy relies on the monitoring capabilities of the transport equipment and the ability of the network management software (NMS, SDN) to proactively notify the network operator of any potential problems. A good example of this is the per channel power monitoring capabilities enabled by embedded Optical per Channel Monitoring (OCM) technology in the DWDM hardware. Not only can this capability be used to monitor the power level of each wavelength at different points across the network, but with threshold crossing alarms in the network management system able to notify the network operator that the attenuation between two monitoring points has increased, the network operator can then take preventative action before the lambda is impacted. Another example is the active monitoring of the protection path so that any failures in the protection path can be identified and repaired ensuring that it will be available in the event it is needed.

Looking further ahead, a key use case for SDN is SLA-aware service assurance with the network software monitoring the SLA of each service and then proactively taking steps to ensure the SLA is maintained in the event of any degradation. Examples of actions could include rerouting the traffic, reprioritizing the traffic, assigning more bandwidth, and tearing down or rerouting lower priority services.

REDUCING HUMAN ERROR

Human errors can occur at any stage including installation, network turn up, and service provisioning. Strategies to reduce the risk of human errors during installation include shipping fully configured and tested nodes directly from the factory, integrated system-on-a-blade solutions with minimal cabling, and Zero Touch Commissioning. Zero Touch Commissioning works to simplify the installer's job by downloading the software and configuration automatically from a central server based on the location of the node in the network.

Integrating Optical Time Domain Reflectometer (OTDR) with Raman amplifiers can ensure that the Raman amplifiers will not be impacted by excessive reflections from splices before they are turned on. During network deployment, DWDM equipment should be able to auto-balance with the amplifiers automatically adjusting their power levels based on the span loss and number of channels, even if no channels are present. As DWDM channels are added or torn down, the network should auto-balance both the amplifiers and the individual channels in real time.

Additional approaches to reducing human error include point-and-click provisioning in the NMS and tight integration between planning tools and the NMS, control plane, and SDN.

REDUCING THE TIME TO REPAIR FAULTS

While it may be impossible to eliminate service impacting faults entirely, the ability to quickly locate and repair faults can, by reducing the Mean Time to Repair (MTTR), have a significant, positive impact on network availability.

A number of tools are available at different layers of the network to identify the location of faults. For the optical layer, these tools include per channel power monitoring and integrated OTDR. Per channel power monitoring when embedded in nodes throughout the network can help to quickly locate impacted spans, while an integrated OTDR can pinpoint the exact location of any fiber cut. OTN Tandem Connection Monitoring (TCM) can be used to quickly identify the domain where a failure has occurred. Y.1731 and 802.1ag CFM can be used to identify where in an Ethernet network any fault has occurred.

A critical factor impacting MTTR on faults concerning transport network equipment is the availability of spares. Universal switching, programmable and tunable interfaces, and ROADM-on-a-blade/amplifiers that support a wide range of span losses, can all reduce the number and types of spares that must be stocked. Furthermore, protection and restoration can also be used as strategies for minimizing the impact of long repair times.

DEALING WITH FIBER DEGRADATION

As the quality of the fiber degrades over time, services could be impacted by decreases in OSNR and increases in bit error rates. A number of tools are available to minimize this risk. Accurate fiber characterization is critical prior to network planning to ensure the network is planned with correct information.

Additionally, a sophisticated optical path computation engine is required to accurately plan the network to ensure that there is sufficient margin for end of life. Furthermore, integrated PRBS test and loopback should be used to soak test any new wavelengths to ensure that the performance is as predicted. Finally, especially in long haul networks, it is critical to have link control software that can make the necessary adjustments to power levels to ensure robust performance against changes in attenuation or other fiber parameters.

STRATEGIES FOR COST-EFFECTIVE HIGH AVAILABILITY IN THE TRANSPORT LAYER

As described previously in this white paper, there are a wide variety of features and options that can be used to deliver high availability. The challenge is to pick the right mix to deliver the appropriate level of availability at the lowest cost. The correct strategy will depend on a variety of factors including:

- Risk of fiber cuts or other non-equipment related failures
- Content of any Service Level Agreements
- Mix of services and the required availability for these services
- Skill and knowledge base of network operations staff
- Role of the transport network in the multi-layer strategy for high availability

For example, an operator wanting to deliver business service with a differentiating SLA in a market with a higher risk of fiber cuts will need a different approach than an operator offering residential services in a market with a low risk of fiber cuts. However, the following suggestions may be useful when determining the appropriate solution:

1. Look for solutions that can cost effectively deliver 50 ms 1+1 protection for coherent channels.

Solutions for 1+1 protection including 1+1 OCh and 1+1 OMS that avoid the need to duplicate expensive coherent line interfaces can provide a highly cost-effective option for surviving fiber cuts and other failures. However, as described previously, delivering 50 ms fail-over times for coherent signals is highly challenging due to the need to retune the DSP and in some scenarios detect the failure even though the receiver is seeing multiple channels and cannot rely on Loss of Signal. Fortunately, after significant Coriant R&D investment, our long haul solutions can deliver 1+1 OCh line protection and 1+1 OMS protection with 50 ms fail-over even with coherent colorless add/drop. Furthermore, these protection mechanisms monitor the health of the protection resources and provide options for reversion so that the original working path can be reused once it has been repaired.

2. Look for solutions that can deliver a wide range of protection and restoration options.

Coriant metro and long haul transport solutions that support a wide range of protection and restoration schemes enable operators to tailor the resiliency mechanism to the desired availability for each individual connection or service. Protection options include client facility protection based on Link Aggregation and APS/MSP 1+1 and client side network protection including Y-cable. Line side protection options include 1+1 OCh at the optical layer, SNC for OTN and SONET/SDH, G.8031 VLAN protection and G.8032 ERP for Carrier Ethernet, and 1:1 MPLS-TP protection. 1+1 OMS and 1+1 OTS are also supported. Key features of the common ASON/GMPLS control plane include dynamic restoration at the optical and electrical layers and the ability to combine protection and restoration.

3. Look for a solution that cost effectively delivers high node availability.

High node availability is a cost-effective insurance policy against a single point of failure taking out all the traffic passing through a node. Features including redundant power, redundant fans, redundant shelf controllers, and distributed or redundant fabrics are available across the Coriant packet optical transport portfolio. Examples of ways in which Coriant has made high node availability cost-effective include fabricless switching in the Coriant® 7100 Nano™ Packet Optical Transport Platform and Coriant® 7100 Pico™ Packet Optical Transport Platform, 1:5 redundant fabrics in the Coriant® mTera® Universal Transport Platform (UTP) where a single backup fabric is able to protect five active fabrics, and the processors on the port cards of the 7100 Pico that run the system software. Coriant optical transport products will also continue to forward traffic even if the system control function has failed.

4. Look for a solution that delivers low cost features to reduce human error, speed repair times, and tolerate fiber degradation.

Features to reduce human error, enable proactive fault prevention, speed repair times, and tolerate fiber degradation can largely be delivered with software or low incremental hardware cost and offer a very high return on investment.

Examples of key features within the Coriant packet optical transport portfolio that can reduce human error include integrated ROADM-on-a-blade in the 7100 Series and mTera® UTP that minimizes the risk of cabling errors. Coriant long haul solutions support Rack and Stack where the rack will be configured in the factory based on planning tool files and shipped as is to speed installation and reduce the risk of human error. Additional features to reduce human error include auto-balancing of amplifiers and per channel power levels. This is supported even in low cost fixed WDM configurations based on the Coriant® Pluggable Optical Layer thanks to cost-effective automatic gain management based on power tone. Human error is also minimized by Coriant® Transport Network Management System (TNMS) point-and-click provisioning and tight integration between the planning tools/path computation engines and TNMS. Another approach to minimizing the risk of human error is to leverage Coriant network deployment services.

Proactive downtime prevention and fast repair times are enabled by extensive performance monitoring including per channel power monitoring offered across the optical portfolio including the Pluggable Optical Layer, which supports a pluggable OCM. Extensive performance monitoring and fault management is supported at multiple layers including OTN, Carrier Ethernet with Y.1731/802.1ag, and MPLS-TP. Coriant optical solutions also support integrated OTDR for quickly and accurately locating fiber cuts.

As discussed previously, sparing can also impact the time to repair. Coriant solutions simplify sparing, as a wide range of services can be supported on a minimal set of port cards, including the OSM universal switching cards for the mTera® UTP, which support OTN, Carrier Ethernet, and MPLS-TP as defined in software on an individual port or virtual interface basis. Coriant can also offer managed spare services.

Additionally, a key differentiator for the Coriant long haul solution is the link control software described in the Coriant white paper Maximizing 100G+ Reach in Long Haul Networks with Challenging Fiber Conditions. This link control is able to tolerate large increases in attenuation on one or more spans by adjusting the power levels of each wavelength and the gain of each amplifier.

5. Do not cut corners on upfront planning.

Accurate network planning is a key requirement for maximizing network availability in optical transport networks that is often overlooked. The starting point for this activity is accurate fiber characterization, which is a service that can be provided by Coriant and its partners. Sophisticated planning tools can then take this information and accurately model the optical layer ensuring sufficient margin for end of life. Coriant can also work with you to analyze the topology of your network to model the current risk of downtime and suggest optimal changes to cost effectively increase network availability by modifying the topology, network equipment, resiliency mechanisms, and/or protocol settings.

SUMMARY

While the need for high availability has never been greater, a wide variety of tools are available to help prevent nodes from failing, prevent failures from impacting service availability, reduce human error, and speed the time to repair. The challenge is to implement the right tools to cost effectively maximize network availability considering the risks to the network and the requirements of specific services.

Coriant metro and long haul packet optical transport solutions offer a wide range of capabilities to enable network operators to achieve these goals including protection and restoration features such as sub-50 ms 1+1 OCh and 1+1 OMS protection for coherent channels, cost-effective solutions to eliminate common equipment single points of failure, and a variety of tools for proactively preventing downtime and shortening repair times.

ABOUT CORIANT

Coriant delivers innovative, dynamic networking solutions for a fast-changing and cloud-centric business world. The Coriant portfolio of SDN-enabled, edge-to-core transport solutions enables network operators to reduce operational complexity, improve utilization of multi-layer network resources, and create new revenue opportunities. Coriant serves leading network operators around the world, including mobile and fixed line service providers, cloud and data center operators, content providers, cable MSOs, large enterprises, government agencies, financial institutions, and utility companies. With a distinguished heritage of technology innovation and service excellence, forged by over 35 years of experience and expertise in Tier 1 carrier networks, Coriant is helping its global customers maximize the value of their network infrastructure as demand for bandwidth explodes and the communications needs of businesses and consumers continue to evolve. Learn more at www.coriant.com.

These trademarks are owned by Coriant or its affiliates: Coriant®, Coriant Dynamic Optical Cloud™, Coriant Transcend™, Coriant CloudWave™, mTera® , Nano™, and Pico™. Other trademarks are the property of their respective owners. Statements herein may contain projections regarding future products, features, or technology and resulting commercial or technical benefits, which may or may not occur. This publication does not constitute legal obligation to deliver any material, code, or functionality. This document does not modify or supplement any product specifications or warranties. Copyright © 2018 Coriant. All Rights Reserved. 74C.0004 Rev. B 01/18