



ADVA FSP 3000 CloudConnect™

Terascale Transport for the Cloud of Tomorrow



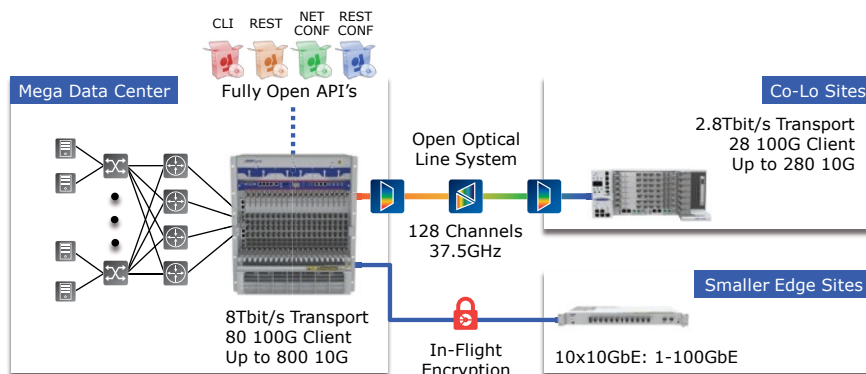
The Rise of DCI Networks

Consumer habits, business practices and bandwidth consumption trends are ever changing, and so must the networks that serve them. The two biggest trends in bandwidth consumption have been a shift to mobile access and cloud service. This has led to the rise of successful Internet Content Providers (ICPs) and Data Center Operators (DCOs). And with their rise has come an eagerness to question the way their underlying networks are architected and built. No longer satisfied to reuse telecommunications equipment designed for the internet of old, they have spearheaded efforts to create Data Center Interconnect (DCI) networks that better meet their needs.

From Telecom to Datacom

The very fundamental architecture of how data is transported is undergoing a sea of change. Where content is located, and how it is delivered is changing. And we purveyors of optical transport must change with it, or risk being obsoleted.

To understand where networks are headed, one must first understand where they have been with telecommunication. These tended to be redundant, protection switched physical rings with add/drop nodes at each city passed. Packets were encapsulated into OTN containers, with OTN switches used for finer granularity add/drop. The shipping lanes on fiber paths were controlled via a network controller with the goal of optimizing said paths, with little to no optimization of what might be contained within the containers. Physical transport and data transport were isolated entities. While routers used various protocols to try their best to optimize packet throughput, they were at the mercy of what the OTN network could and couldn't do. As bandwidth consumption became more mobile and real-time, the only way IP/OTN networks could keep up was with heavy overprovisioning beyond all reason. And overprovisioning leads to waste, the bane of Data Centers' existence.



A Better Way

Seeing the shortcomings in their early networks, ICPs and DCOs developed a better way. Data distributed across a load-balanced cloud is already redundant many times over, so out went fiber level protection switching, including the redundant fiber rings required. The switching granularity of optical wavelengths and OTN containers is too coarse for their needs, so out went wavelength and OTN switching. There was little value in trying to optimize fiber, OTN, and transport paths, as what was really needed was optimization of end-to-end multilayer flows, so out went any intelligent controllers. Instead, everything is tied to a common orchestration layer.

Times have changed. The enemy is no longer longevity of equipment, with a focus on how many decades the equipment can be milked for revenue. The enemy is now technology obsolescence, with a focus on how quickly can equipment be installed and removed, while in-service. The emphasis is no longer minimizing truck rolls and the cost they entail, but rather minimizing power consumption and footprint. The competition is no longer who can have the most stalwart network, but who has the newest, most bleeding-edge network.

How DCI is Different

Yes, a new age has dawned in networking. All the world's information resides in mega data centers around the globe, in massive leaf-and-spline horizontal arrays of servers. These data centers themselves are interconnected with massive DCIs, both between themselves, and to local content caching sites. The fiber optic links are encrypted in-flight to protect the value of the information being carried. The equipment may be at owned or leased sites, private or shared. The optical transport gear and the associated network elements are fully open, both the hardware and software. Rather than make decisions best for local conditions, the transport network now defers to an orchestration layer with full visibility of multilayer packet flows and asset utilization.

How the Network has Changed

- The intelligence to optimize traffic no longer resides in network layer
- Traffic has shifted from packets to flows
- Horizontal traffic now dominates vertical
- Global caching makes encryption everywhere necessary
- Closed network management systems no longer tolerated

The ADVA FSP 3000 CloudConnect™ Solution

ICPs' and DCOs' needs are unique, so why should they settle for anything less than DCI optimized optical transport gear? They need something small and efficient today, yet will scale indefinitely for tomorrow. They have to fit into smaller shared co-location facilities, yet they would like the same equipment to form the backbone of their network. They want to minimize the number of inventoried items, yet still maintain flexibility to support new client interfaces as they arise. They want their network managers to have full visibility of the network, including the transport gear, yet they don't want to have to deal with archaic physical network managers. In other words, they need the ADVA FSP 3000 CloudConnect™.

DCI's Sweet Spot

With all the hype around 1RU monolithic "pizza box" style transport products, one would think that smaller must be better. But ADVA Optical Networking did not assume anything. We worked in close association with Internet Content Providers (ICPs) and Data Center Operators (DCOs) and asked them what was best. Their answer? Something around 4RU.

Turns out there are problems with 1RU and monolithic products. 1RU products require too many chassis to manage, as for each chassis, there are redundant power supplies, power connections, fan modules, and management interfaces to deal with. They are a great place to start a transport service, but once bandwidth demand outgrows the chassis capacity, scaling becomes awkward. Making an undersized chassis monolithic (fixed ports) only compounds the problem, as now a different chassis variant must be stocked for every port type. And as new client types are introduced, existing chassis must be ripped out and replaced with new ones.

The ADVA FSP 3000 CloudConnect™ was designed by and for DCI customers. The 4RU chassis is in the "Goldilocks zone", as it is large enough to both scale and reduce inventory sprawl, but still small enough to squeeze into available rack space while minimizing co-location leasing costs. Where needed, it functions as a monolithic chassis, with the majority of common equipment integrated; but also where needed, it acts as a blade chassis, with redundant, removable shelf controllers, data ports on line cards, and rear ac-



cessible redundant power supplies and connections. In this case, multiple personalities equals optimal configuration. And for DCI customers who have committed to blade server architectures and want a similar transport solution, the FSP 3000 is also available in an ETSI 300mm deep 12RU chassis.

FSP 3000 QuadFlex™ Card

The DCI application requires unprecedented levels of flexibility, and the FSP 3000 QuadFlex™ line card delivers. The single slot line card offers an astounding 800Gbit/s of capacity, shared among 4 client ports and 2 network ports, while drawing only typically 140 Watts. The 4 client ports are each 100Gbit/s QSFP28 interfaces supporting 100Gbit/s Ethernet and OTU4 protocols, as well as SR4, LR4, CWDM4, and AOC reaches.

The dual network ports can be configured a variety of ways to optimize distance vs bandwidth tradeoff, as well as grow incrementally from 100Gbit/s to 400Gbit/s transport capacity. At the lower end, a single QSFP28 may be populated and a single network port used for 100Gbit/s transport over 3500 kilometers! This capacity may be doubled with no reach penalty. At the high end, all 4 QSFP28 clients are used with both network ports to create a pair of 200Gbit/s channels that can reach up to 800 kilometers. Unique to the ADVA FSP 3000 CloudConnect™ solution, the line card may also be configured for an in-between 300Gbit/s transport capacity. When thus configured, 3 QSFP28 clients are used in conjunction with a pair of network ports running 150Gbit/s per channel, for distances up to 1500 kilometers.



A Complete Solution, OLS Included

What good is having a small 1RU transponder chassis, if it has to connect to a huge Optical Line Shelf (OLS)? The ADVA FSP 3000 CloudConnect™ includes channel multiplexers, optical amplifiers, all the equipment necessary for end-to-end transport in most DCI applications. In fact, a key part of the solution is a 1RU channel multiplexer that will grow from a single 100Gbit/s channel all the way to 128x200Gbit/s channels, without the need for unsightly rats' nests of cabling.



Another important part of the ADVA FSP 3000 CloudConnect™ solution is a single slot width EDFA optical amplifier that greatly simplifies the total DCI solution. All that is needed is simple cabling between the mux and amplifier. The output of the amplifier is ready for your fiber pair.

The ADVA FSP 3000 CloudConnect™ Solution

- Best-in-class metrics: 1.4Tbit/s per Rack Unit; 2Tbit/s per kW
- Multiple chassis configurations available, including 4RU and 12RU
- Same line cards may be used across all chassis configurations
- Scales from a single 100Gbit/s to 128x200Gbit/s
- Chassis configurations include optical line shelf components



The Price, Power, Size They Need Today with the Scale They Will Need Tomorrow

While the word "Terabits" gets thrown around a lot in DCI these days, the truth is that Terabits is more of an end goal. For many ICPs and DCOs, today's requirements can be as little as 100Gbit/s. And just because ICPs and DCOs are growing like crazy does not mean they are willing to waste money, power, or space. The price, power, and space margins are razor thin in DCI, and an optical transport solution must deliver from the first Gigabit to the last Terabit. The ADVA FSP 3000 CloudConnect™ was designed by ICPs and DCOs to meet their needs of today, while being able to scale seamlessly for their needs of tomorrow.

From that First 100Gbit/s Port

The minimum system configuration of the ADVA FSP 3000 CloudConnect™ is a single chassis with a single FSP 3000 QuadFlex™ line card. Since the line card will support a single 100Gbit/s QSFP28 client feeding into a single 100Gbit/s network port, that is the minimum configuration. For customers who wish to start out with a 400Gbit/s network channel plan, the line card may be fully configured for maximum capacity and throughput, and pay-as-you-grow pricing plans are available, making the transition to 400Gbit/s networks as painless as possible. For customers who have the opposite problem and are not ready for 100Gbit/s granularity, 10Gbit/s multiplexing line cards are available.



This minimal configuration is able to grow from that first 100Gbit/s port to a total of 28x100Gbit/s client ports. This is sufficient bandwidth to fill up 14x200Gbit/s WDM channels. An interesting midway point is 3 line cards and an amplifier, providing 1.2Tbit/s of optical transport, enough to join the Terabit club.




Filling an Existing C-Band System

The majority of WDM transport networks today operate over a 40 channel plan in the C-band on 100 GHz grid spacing. For many customers, they want to move to the latest DCI optical transport technology, but they want to operate over their existing 40 channel grid. For this customer set, the ADVA FSP 3000 CloudConnect™ can be configured as 3 stacked 4RU chassis with a 40CSM channel multiplexer and single amplifier line card. This configuration provides 8Tbit/s of optical transport over an existing 40-channel WDM grid.

Scaling to 25.6Tbit/s over a Single Fiber Pair

But some ICPs and DCOs are building global scale computing clouds. For these customers, they can never have enough transport bandwidth between data centers. The performance of their compute cloud is directly proportional to the inter-connection bandwidth, and they could use 100 Terabits today if available. For these customers, the same ADVA FSP 3000 CloudConnect™ solution that begins at a single 100Gbit/s port can scale all the way to 256x100G clients, for total of 25.6Tbit/s transported over a single fiber pair. What's more, the 25.6Tbit/s fits into a single 19" rack, including channel multiplexers and amplifiers.



FSP 3000 Cloud Connect™: Unprecedented Scalability

| 100G Clients | QuadFlex™ Line Cards | WDM Channels | Capacity (Tbit/s) | Rack Units |
|--------------|----------------------|--------------|-------------------|------------|
| 1 | 1 | 1 | 0.1 | 4 |
| 4 | 1 | 2 | 0.4 | 4 |
| 12 | 3 | 6 | 1.2 | 4 |
| 28 | 7 | 14 | 2.8 | 4 |
| 80 | 20 | 40 | 8 | 14 |
| 256 | 64 | 128 | 25.6 | 40 |

DCI Software Needs Are Unique

Data Centers have state-of-the-art management software. Server loads are monitored and optimized, with loads continually being balanced across all available resources within their walls. However, what happens when the same is tried between Data Centers, outside their walls? A roadblock is hit. Traditional network management software is not compatible with Data Center software, thus preventing the latter from truly optimizing end-to-end multilayer packet flows. DCI optimized optical transport such as the ADVA FSP 3000 CloudConnect™ removes that roadblock, allowing the optical transport piece to fit effortlessly and seamlessly into existing management systems, allowing global optimization of resource allocation.

Direct APIs to NE for Simple Connectivity

The level of sophistication required by DCI customers varies widely. For simple connectivity of point-to-point networks, often direct APIs are the right choice. In this case, network elements are provisioned using Command Line Interfaces (CLI), sometimes with scripting routines. One advantage of CLI is that it is human readable text, as opposed to encoded SNMP. The ADVA FSP 3000 CloudConnect™ CLI is similar to popular router interfaces. If someone can provision a router, then they will have no trouble provisioning Terabit transport networks. Basic Representational State Transfer (REST) interfaces can be used here, as well.

Network-level APIs for Meshed Networks

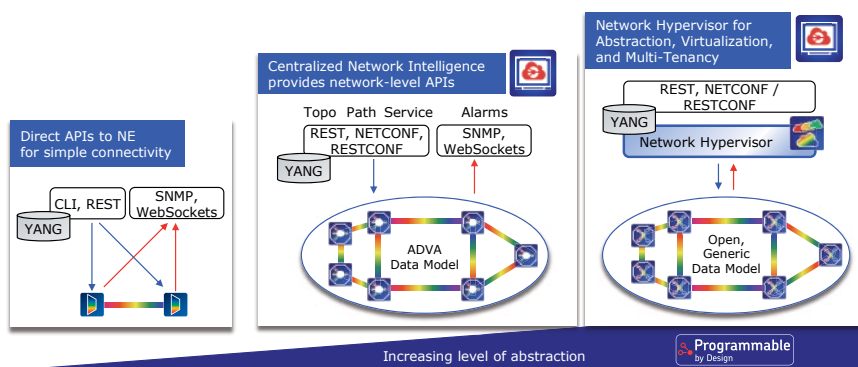
When networks move beyond simple point-to-point links, and into more complex meshed network elements, then network-level APIs are required. The aforementioned RESTful interface is one example. Another is Network Configuration Protocol (NETCONF) and uses an Extensible Markup Language (XML) to communicate to network elements. NETCONF is already popular with routers and switches. RESTCONF attempts to combine the best attributes of both REST and NETCONF, by using a RESTful interface running over HTTP for accessing data stored in NETCONF.

Network Hypervisor for Abstraction, Virtualization and Multi-Tenancy

When full optimization of end-to-end multilayer data flows is desired, then a greater level of intelligent control is required. The DCI transport network must become part of the overall resources being optimized, but this is prevented by the complexity inherent in optical transport. The solution is to add an additional hypervisor, a transport network hypervisor that abstracts the entire transport network and presents it to the orchestration layer in a format that it can understand. The ADVA FSP 3000 CloudConnect™ solution includes a transport network hypervisor with fully open Northbound interfaces. It can control anything, or be controlled by anything. It is not a network controller itself, but rather a translator that lets anyone or anything query and control the transport network. It sounds straightforward, but the end result is very powerful. For example, now an orchestration application attempting to load balance across virtual machines at different geographically dispersed locations can include the transport network in its optimization. No longer are higher layer applications held hostage to the transport network.

Another example of the power of abstracting the transport network with a network hypervisor is multi-tenancy. A typical Service Provider will use a common transport network to serve multiple clients. Currently, each of these clients may have bandwidth connectivity, but they must go through the Service Provider for all network management. A multi-tenancy capable network hypervisor allows each client to have their own virtual transport network, based upon the real physical assets available. Each client sees their own and only their own portion of the network and can configure it within pre-determined configuration rules. The Service Provider maintains visibility of the entire network. This greatly facilitates a Service Provider selling as many services over a single infrastructure as possible. Customers

who would have demanded their own manageable network can now be sold a shared network.



Fully Open and Programmable APIs

- Transport SDN is gradually turning from 'OpenFlow Hype' towards Open Programmable Networks based on YANG models
- Customized APIs depending on customer requirements
- Direct APIs (CLI, REST) at NE level for point-to-point network
- Network level APIs (Topology, Service, Path Computation) for simple mesh networks
- Network Hypervisor for Network Abstraction, Virtualization, and Multi-Tenancy

For More Information

ADVA Optical Networking SE
Campus Martinsried
Fraunhoferstrasse 9a
82152 Martinsried/Munich
Germany

ADVA Optical Networking North America, Inc.
5755 Peachtree Industrial Blvd.
Norcross, Georgia 30092
USA

ADVA Optical Networking Singapore Pte. Ltd.
25 International Business Park
#05-106 German Centre
Singapore 609916

info@advaoptical.com
www.advaoptical.com

About ADVA Optical Networking

At ADVA Optical Networking we're creating new opportunities for tomorrow's networks, a new vision for a connected world. Our intelligent telecommunications hardware, software and services have been deployed by several hundred service providers and thousands of enterprises. Over the past twenty years, our innovative connectivity solutions have helped to drive our customers' networks forward, helped to drive their businesses to new levels of success. We forge close working relationships with all our customers. As your trusted partner we ensure that we're always ready to exceed your networking expectations. For more information on our products and our team, please visit us at: www.advaoptical.com.



Optical+Ethernet Innovation • Speed for Customers • Trusted Partner

