Technology Brief

Decoupling Aggregation and Transponding

A flexible, efficient, operationally simple Data Center Interconnect architecture

Data Center Interconnect (DCI) traffic comprises various mixes of 10 GbE and 100 GbE client-side services on each link. In a typical application, DWDM is used to maximize the quantity of traffic that can be carried on a single fiber. Virtually all available products for this function combine aggregation and transponding into a single platform. These platforms aggregate multiple 10 GbE services into a single 100G wideband signal and then transpond that signal onto a wavelength for multiplexing onto a single fiber along with other wavelengths. This design frequently results in wasted rack space and unused, stranded bandwidth.

A New Approach to Aggregation and Transponding

Decoupling aggregation and transponding into separate platforms is a new design approach made possible by the unique feature set of the modular Fujitsu 1FINITY™ Transport family, specifically the T100 and T400 blades. This approach employs these two blades in place of the typical converged platform, separating the two halves of the formerly converged aggregation-transponding function:

1. The 1FINITY T400 serves as the 10 GbE-to-100G aggregation device.
2. The 1FINITY T100 serves as a separate hyper-dense transponder.

The converged approach and the decoupled approach are depicted together for comparison in Figure 1. Both approaches accommodate 10 GbE and 100 GbE client-side services and both use 200G wavelengths.

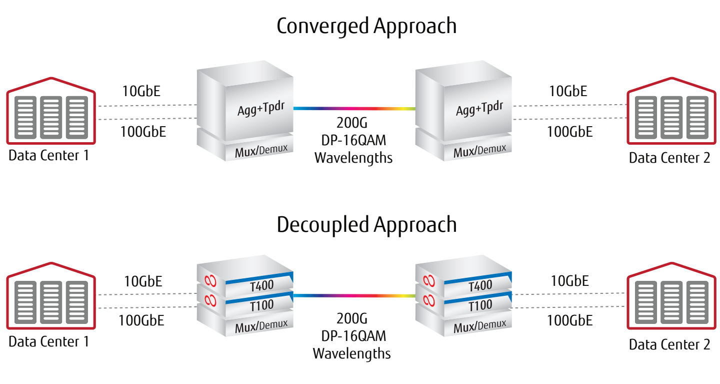

Figure 1: Comparison of Converged and Decoupled DCI Architecture

The Value of Decoupling Aggregation and Transponding Functions

Paradoxically, this decoupled solution may seem less desirable than the typical converged design at first glance. However, it offers several advantages—particularly in DCI applications. These advantages are essentially an extension of the overall disaggregated, blade-centric approach to DCI architecture exemplified by Fujitsu 1FINITY optical networking platforms.

Efficient Use of Rack Space

Physical separation of aggregation and transponding splits a single larger unit into two smaller ones: a dedicated aggregator and a dedicated transponder. As a result, it is easier to find space for both and use up scattered empty 1RU slots, which helps make the fullest possible use of costly physical facilities.

Reducing Stranded Bandwidth

Many suppliers use QSFP+ transponders, which offer programmable 40G or 100G service. Bandwidth can be wasted when aggregating 10 GbE services into 40G, because mathematically 40 is not a factor of 100. This necessitates deployment in multiples of 200G in order to make the numbers work out evenly. The 1FINITY T400 aggregator deploys in chunks of 100G, which keeps stranded bandwidth to a minimum.

Simplified Operations

Operational simplification occurs for two reasons. First, when upgrading to a higher line rate, you simply change out the transponder without affecting the aggregator. Second, when changing the 10 GbE/100 GbE client mix, significantly fewer operational activities are needed in comparison to converged aggregation/transponder platforms. With fewer technician activities, there are fewer errors and therefore less cost.

Comparing Decoupled and Converged Approaches

A Side-By-Side Comparison of Decoupled and Converged Designs

The areas of improved value using the decoupled 1FINITY design can be highlighted by comparing them directly with the same areas under the typical converged aggregation/transponding approach. Fujitsu engineers performed a series of comparison tests to provide concrete metrics in each area. Functional elements used to make the comparisons that follow are shown opposite.

Reduced Rack Space Usage

Consider a data center from which:

- 4 × 200G worth of traffic is routed to five different data centers.
- Each of these demands comprise 6 × 100 GbE and 20 × 10 GbE client-side inputs.

Converged: This scenario can be met with five converged aggregation/transponder platforms, one for each of the five data center connections, based on the functional elements of the comparison. These five converged 2RU platforms will require total rack space of 10RU (excluding the multiplexer).

Decoupled: With the decoupled design, you would need five T100 blades and a single T400 for aggregation. This brings the total rack space need down to just 6RU (again, excluding the multiplexer). This 40% rack space reduction is shown in Figure 2.

Functional Elements of the Comparison

1FINITY T400

1FINITY T100

1FINITY T400: 1 Tbps Layer 1 aggregator

- The T400 is a 1RU blade that supports 100 × 10 GbE ports (via QSPF+ ports with 4:1 breakout cables) aggregated into 10 × 100GBASE-R QSFP28 ports

1FINITY T100: 800G DCI transponder

- The T100 is a 1RU blade that supports 8 × 100 GbE QSFP clients and 4 × 200 GbE CFP2-ACO narrowband line ports.

Note: These two devices support CR4 copper connections between the T400 100G port and the T100 100G client port.

Typical converged platform

- This device is typically a 2RU chassis with 20 QSFP+ on the client side, that can be configured as 10 GbE or 100 GbE and 4 × 200G narrowband line-side ports.

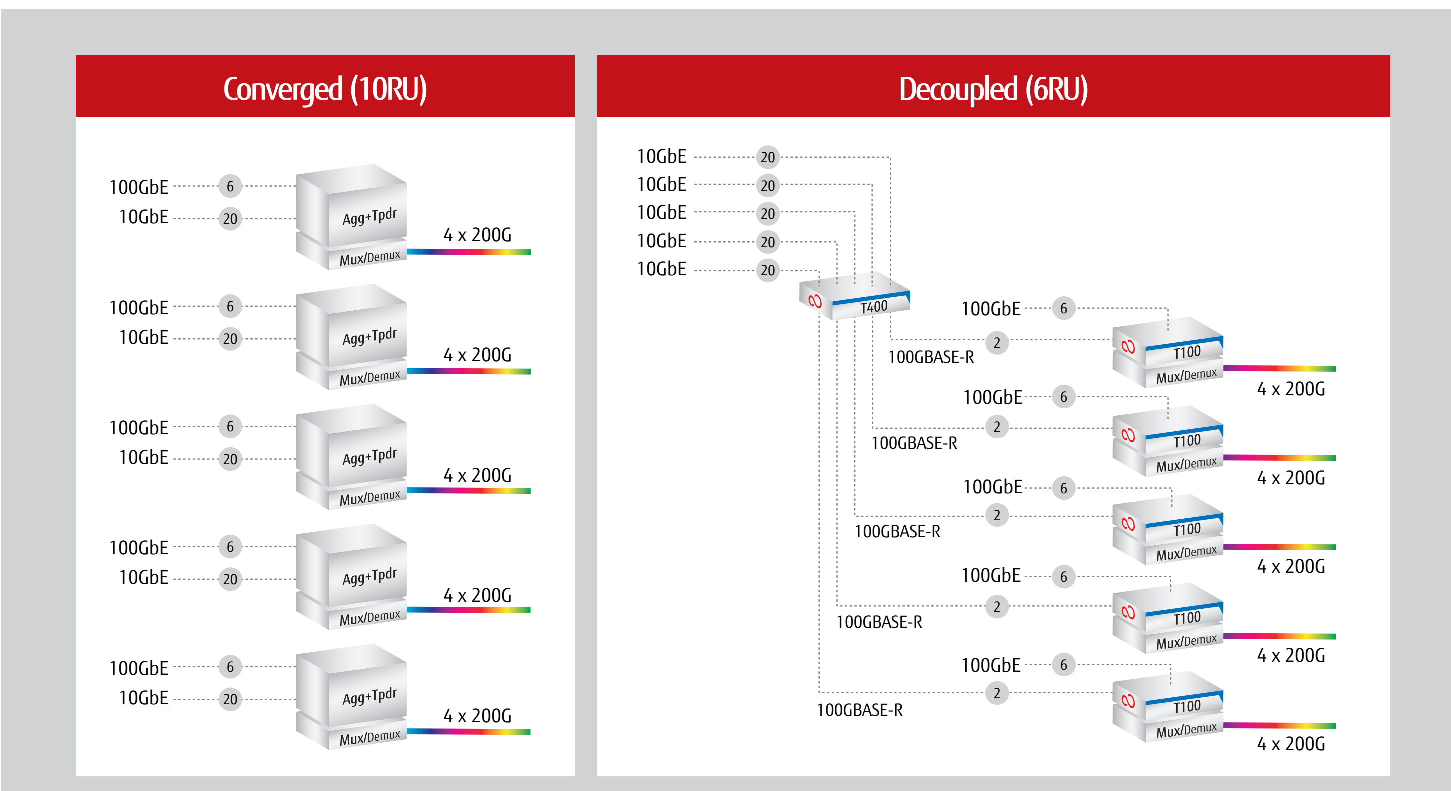

Figure 2: Space savings (10RU versus 6RU)

Bandwidth and Rack Space Efficiency

Less Stranded Bandwidth

To illustrate the problem of stranded bandwidth, it is necessary to focus on port usage and the specific service mix in each scenario. Essentially, issues with stranded bandwidth only arise whenever there is a need to provision an odd number of 10 × 10 GbE blocks or 100 GbE services.

This comparison assumes the following service mix between two sites:

- 10 × 10 GbE
- 7 × 100 GbE

Converged: With the converged design, client ports are QSFP+ or QSFP28. Since these two pluggables have the same form-factor, the cage can accommodate either.

- For 10 GbE, the port is configured for QSFP+
- For 100 GbE, the port is configured for QSFP28
- QSFP+ ports are used on the client side and configured for 4 × 10 GbE, or QSFP28 ports are used for 100 GbE

The essence of the problem is that since 40 is mathematically not a factor of 100, blocks of 5 QSFP+ or 200G are provisioned as 20 × 10 GbE or 2 × 200 GbE.

To meet the 10 × 10 GbE and 7 × 100 GbE demands, one block of 200G is configured for 10 GbE and 4 blocks are configured for 100 GbE. This configuration accommodates 20 × 10 GbE and 8 × 100 GbE over 5 wavelengths, resulting in 200G of stranded bandwidth.

Decoupled: With decoupling, the following port assignment is possible:

- 10 × 10 GbE are aggregated into one 100 GbaseR by the 1FINITY T400
- This 100 GbaseR service consumes one of the eight 100 GbE QSFP28 client ports on the 1FINITY T100
- The remaining seven 100 GbE QSFP28 client ports are used for the remaining seven 100 GbE service demands

In sum, with a converged approach, five wavelengths are needed to carry the 10 × 10 GbE and 7 × 100 GbE with 200G of capacity being stranded. However, the versatility of the decoupled approach allows the demand to be met with four lambdas and no stranded bandwidth. This stranded bandwidth reduction is shown in Figure 3.

Simplified Operations when Changing a Transponder

Another area where the increased flexibility of the decoupled approach reduces cost is simplification of operational activities when transponder changes are needed.

Consider the scenario shown in Figure 4, where there are: 20 × 10 GbE and 6 × 100 GbE between two data centers.

Situations occur where a new technology becomes available, such as a 400G transponder. Fujitsu might, for instance, release a 1RU 1FINITY blade that has 8 × 100G clients and 2 × 400G wavelengths. A transponder upgrade in a 2RU converged chassis is also possible.

A theoretical comparison of the fiber connection and provisioning steps for changing the transponder under each paradigm is shown in Table 1. With the decoupled approach, a technician manipulates fewer connections and issues fewer commands in total, thus making the whole procedure operationally simpler. Inevitably, performing fewer, simpler activities reduces the likelihood of mistakes.

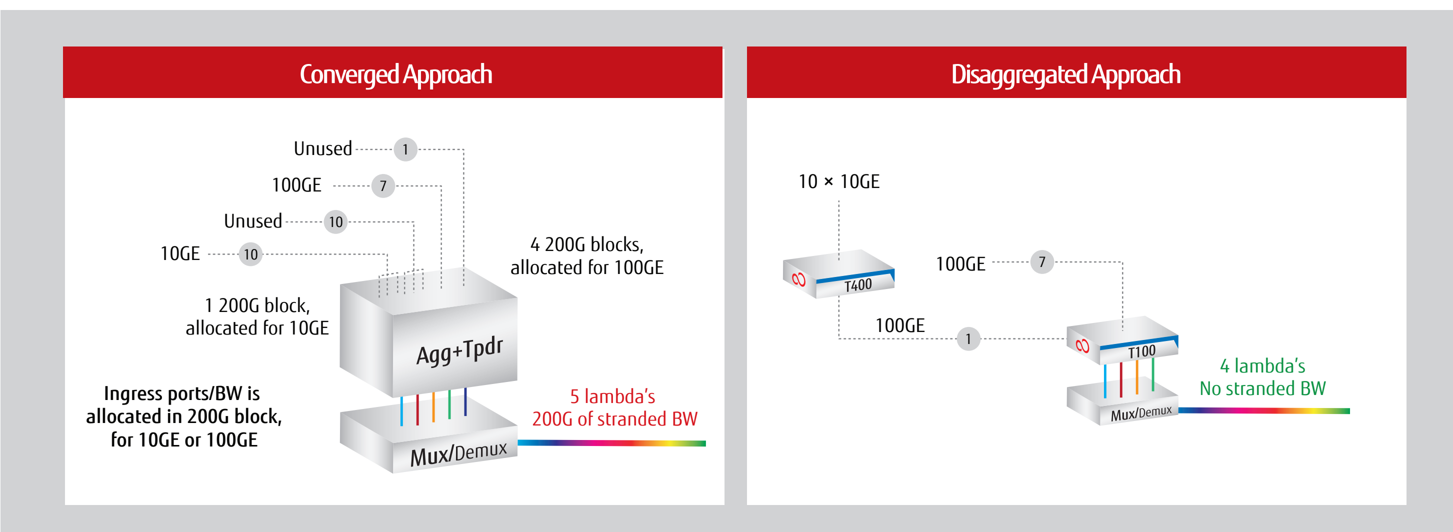

Figure 3: Less stranded bandwidth (200G versus none)

Operational Simplicity

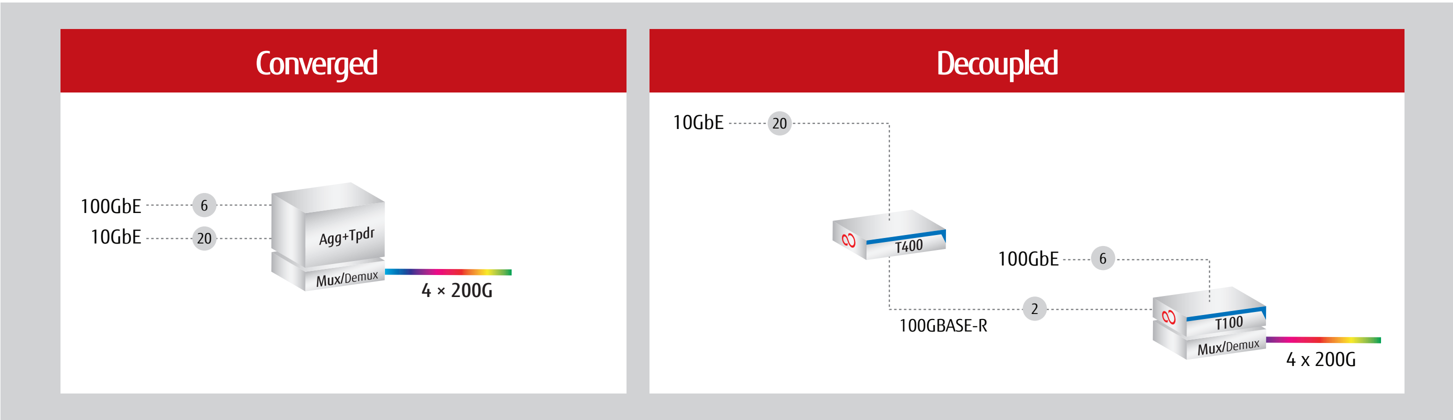

Figure 4: Simpler operation (fewer steps)

Converged Platform	Decoupled 1FINITY Platform
Disconnect 11 connections (6 × 100 GbE and 5 × 40G QSFP+)	Disconnect 8 connections (6 ×100 GbE and 2 ×100 GBASE-R)
Install new 2 × 400G converged transponder	Install new 2 × 400G transponder
Reconnect 11 connections	Reconnect 8 connections
Configure 1 block of 200G as QSFP+	
Configure 5 QSFP+ as 4 × 10 GbE	
Configure 3 blocks of 200G as QSFP28	
Configure 6 100 GbE clients	
Provision 8 transponders	Provision 8 transponders
Provision 2 400G wavelengths	Provision 2 400G wavelengths

Table 1: Comparison of the operational steps for changing transponders

Simplified Operations when Changing the Aggregation Mix

The client mix is not static; there is always a trend to move from smaller to larger bandwidth services. With DCI, there is a trend from 10G between data centers to 100G between data centers. Again, the increased flexibility of the decoupled approach reduces cost via simplification of operational activities when changes are needed.

Using the same scenario in Figure 4, let's assume there is a change in demand on the client side from 20 × 10 GbE and 6 × 100 GbE to 8 × 100 GbE. A theoretical comparison of the fiber connection and provisioning steps for changing the transponder under each paradigm is shown in Table 2. Again, the procedure is simpler under the decoupled design.

Converged Platform	Decoupled 1FINITY Platform
Disconnect 5 40G QSFP+ connections	Disconnect 2 connections (100 Gbase-R)
Re-provision this block of ports to 100 GbE	
Connect 2 connections to the 100 GbE ports	Connect these 2 connections to the 100 GbE ports
Provision these 2 transponders	Provision these two transponders

Table 2: Comparison of the operational steps for changing the aggregation mix of 10G and 100G clients

Summary

Typical DCI application design employs a single converged hardware platform for both aggregation and transponding functions. The Fujitsu 1FINITY T100 and T400 blades can be used in an alternative design that decouples these two functions onto separate platfoms.

Separating the aggregator from the transponder offers data center operators several advantages, including improved rack space efficiency, reduced stranded bandwidth, and simplified operations. As a result, the Fujitsu 1FINITY platform provides a simple, flexible, and efficient architecture for DCI.

Fujitsu Network Communications, Inc.
2801 Telecom Parkway, Richardson, TX 75082
Tel: 888.362.7763

us.fujitsu.com/telecom

1.0/04.17