Solution Brief

Data Center Interconnect over Carrier Ethernet

Unlocking the AI Driven DCI Opportunity for CSPs

Data Center Interconnect (DCI) has become critical as enterprises accelerate their digital transformation strategies. Longstanding drivers such as cloud migration, storage replication, disaster recovery, hybrid-cloud architecture, and multisite business applications - continue to fuel the need for high-capacity, low-latency connectivity. Increasingly, AI is emerging as the dominant catalyst reshaping the **scale and trajectory** of DCI. Massive east–west traffic is generated as organizations train models, and operationalize distributed inference across enterprise sites, data centers and clouds.

This convergence of traditional requirements with AI driven workflows is redefining what modern DCI must deliver. Interconnect infrastructures must support extreme scale, assured performance, resilient operations, and future-proof security.

As AI-driven traffic patterns intensify, and the deployment of AI data centers is becoming more geographically diverse, aligned with power distribution constraints and data-sovereignty mandates, DCI is becoming one of the fastest growing revenue opportunities for Communications Service Providers (CSPs). But capturing this opportunity requires more than simply scaling bandwidth. CSPs must deliver **high capacity, SLA-based transport with assured performance, robust resiliency, operational efficiency, and cost-effective scalability**.

These requirements align directly with the strengths of Carrier Ethernet, making it a uniquely compelling foundation for next generation DCI.

RAD

Your Network's Edge®

Solution Brief

Data Center Interconnect over Carrier Ethernet

Why Carrier Ethernet Services Are a Perfect Fit for DCI

Modern enterprise DCI requires high capacity, low latency, and assured transport performance across distributed data centers and clouds. Carrier Ethernet is uniquely suited for this, delivering assured, SLA backed performance with the reliability, security, and operational efficiency needed for both traditional workloads and emerging AI ecosystems. Carrier Ethernet appeals to enterprises seeking minimal operational burden, predictable and a cost-effective solution.

Key Advantages of Carrier Ethernet for DCI

Efficient Bandwidth Utilization

Granular bandwidth provisioning allows enterprises to purchase only the capacity they need unlike traditional OTN.

Statistical Multiplexing

Multiple applications can share the same Ethernet interface using VLANs or virtual circuits, while QoS mechanisms (e.g., traffic prioritization and shaping) dynamically allocate bandwidth and ensure latency-sensitive workloads, such as database replication and AI inference - receive priority.

Network Topology Flexibility

Carrier Ethernet supports point-to-point, point-to-multipoint, or mesh topologies, enabling VLAN and subnet extension across multiple data centers and clouds and allowing Layer-2 connectivity between geographically separated facilities.

Performance

Performance is measured per frame/packet at the Ethernet service layer, aligning directly with how applications and customers experience the network. Unlike optical-layer counters used in optical transport, or tunnel/LSP-based measurements in MPLS networks, Carrier Ethernet enables precise monitoring and assured SLA for latency, jitter, and packet loss at the actual service level.

Scalability

Bandwidth can be rapidly provisioned on demand, with elastic capacity that adapts to dynamic AI, cloud, and enterprise workloads.

Solution Brief
Data Center Interconnect over Carrier Ethernet

Security

Private, isolated Layer-2 connectivity with MACsec wire-speed encryption safeguards sensitive data and AI/model traffic. This approach avoids the complexity and tunnel overhead typically associated with IPsec encryption in MPLS environments, while supporting enterprise requirements for data sovereignty and regulatory compliance.

Cost Efficiency

Enterprises pay only for the bandwidth they need and require minimal on-site equipment, typically just an Ethernet port and a CSP demarcation device (NID/NTU), reducing total cost of ownership compared to DWDM-based services.

RAD's DCI Solution

RAD's DCI solution is powered by the **ETX-2i-400G**, a purpose built 400G Ethernet demarcation and aggregation platform engineered for the ultra-high capacity, low latency, and security demands of modern data center connectivity. It delivers assured performance with advanced traffic management, hierarchical QoS, **line rate MACsec encryption, and a quantum safe security framework**, supporting massive data transfer across data centers, clouds, enterprise sites, and edge locations.

To support high service availability, the ETX-2i-400G incorporates robust load balancing and resiliency mechanisms, ensuring optimal link utilization and fast recovery. The platform also features Zero Touch Provisioning (ZTP) and NETCONF/YANG based programmability, enabling seamless integration into SDN and NaaS automation frameworks. This simplifies deployment, accelerates service turnup, and streamlines lifecycle management across dynamic, AI-driven business environments.

To address emerging long-term security challenges, the ETX-2i-400G supports quantum-safe line-rate encryption, including Post-Quantum Cryptography (PQC) to resist future quantum-enabled attacks, Quantum Key Distribution (QKD) integration for environments requiring quantum-generated keys, and hybrid cryptographic modes that combine classical and quantum-safe mechanisms to ensure continuity and interoperability throughout the transition to post-quantum security.

Solution Brief

Data Center Interconnect over Carrier Ethernet

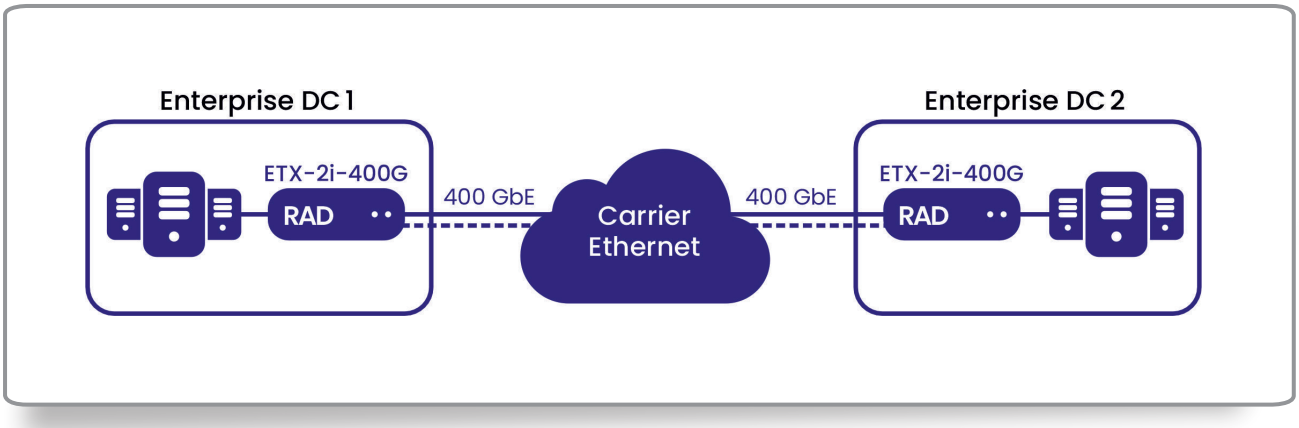

The ETX-2i-400G offers best-in-class traffic management, policing, and HQoS (Hierarchical Quality of Service) capabilities, and it can also handle elephant flows through a "fat-pipe detection" mechanism that identifies high-bandwidth sessions and applies bandwidth policing to prevent them from impacting other traffic.

With optimized bandwidth utilization over WDM or dark fiber, the ETX-2i-400G empowers CSPs to deliver premium, AI-ready DCI services at scale. Its high performance and comprehensive security features make it an ideal foundation for efficient, resilient, and futureproof CSPs DCI offering.

Use Case Example #1:

European CSP Sets a New Standard for DCI Performance and Agility

A leading European CSP selected RAD's ETX-2i-400G to power its carrier grade Ethernet service, delivering high capacity, low latency Data Center Interconnect (DCI) over dedicated dark fiber and WDM optical infrastructure. The solution provides transparent, SLA-backed Layer-2 connectivity ideal for modern data center architectures.

Leveraging the CSP's optical footprint, the service delivers:

- **Bandwidth options up to 400 Gbps** across dedicated fiber or WDM
- **Transparent Layer 2 connectivity**, enabling stretched VLANs, SAN extension, cluster synchronization, storage replication, and AI dataset distribution
- **Low latency and minimal jitter**, essential for synchronous replication and real-time AI inference flows
- **End-to-end service management** with continuous performance reporting

This forms a robust, predictable backbone for active-active data centers, private cloud regions, and **AI-intensive inter-DC applications.**

Solution Brief
Data Center Interconnect over Carrier Ethernet

Service Classes Optimized for All Data Center Traffic

The CSP offers **five differentiated Service Classes** to ensure the right prioritization and QoS treatment across all DCI traffic types:

- **Standard Data** – general enterprise applications
- **Critical Data** – business-critical flows
- **Real Time Data** – ultra sensitive, low-latency traffic (ideal for AI inference and distributed pipeline coordination)
- **Voice Class** – VoIP and unified communications
- **SAN Class** – jitter sensitive storage replication

These classes are enabled directly by the RAD ETX-2i-400G, which provides the advanced HQoS and traffic management capabilities required to enforce the CSP's service policies.

The ETX-2i-400G ensures service integrity with:

- **Hierarchical QoS** for granular per service and per flow prioritization
- **Advanced classification and policing,** mapping VLANs and traffic types into the CSP's QoS framework
- **Strict scheduling, shaping, and congestion avoidance,** preserving latency and jitter guarantees
- **Service aware traffic engineering,** ensuring SAN, real-time, and **AI workload flows** receive guaranteed resources

These capabilities allow the CSP to deliver consistent, predictable performance and uphold stringent SLAs for bandwidth, latency, jitter, and packet delivery.

Real Time Control & On Demand Bandwidth

Real-time agility is enabled by the **programmability of the RAD ETX-2i-400G,** which supports modern **SDN** and **Network-as-a-Service (NaaS)** models. This allows customers to use the CSP's service portal for:

- **Instant bandwidth adjustments,** scaling capacity dynamically to support replication windows, AI bursts, or dataset ingestion
- **On-demand creation of point-to-point Ethernet services,** enabled by SDN-driven provisioning for fast DCI expansion

With secure management and NETCONF/YANG interfaces, and SDN/NaaS readiness, the ETX-2i-400G enables the CSP to automate service activation and deliver a fully elastic, on-demand DCI experience

Solution Brief
Data Center Interconnect over Carrier Ethernet

Reference Use Case Summary

Powered by **RAD's ETX-2i-400G**, this CSP delivers a high-performance Carrier Ethernet based DCI service on its optical network. With rich QoS, real-time bandwidth elasticity, high-capacity transport, and strong SLAs, it provides a scalable, AI-ready, application aware interconnect for modern data centers and distributed enterprise workloads.

Use Case Example #2: DCI for AI-Powered Fraud Detection Platform

A financial institution is deploying a new AI-driven fraud detection and prevention platform to analyze customer transaction data and behavioral patterns across multiple payment types, channels, and banking systems.

To enhance detection accuracy, minimize response times, and strengthen fraud prevention, the company operates three critical locations:

- **Cloud On-Ramp / Azure Gateway facility:** Provides high capacity, secure connectivity to **Microsoft Azure,** where GPU based AI training clusters reside.
- **Primary enterprise data center:** Performs real-time AI inference and scoring on live transactional traffic.
- **Secondary enterprise data center:** Handles model staging, preprocessing, and backup workflows.

These sites must exchange large datasets, real-time inference results, and continuous AI model updates with low latency, strong security, and stringent SLAs to ensure fast and reliable fraud detection.

ETX-2i-400G–Powered DCI Fabric

To meet these demanding interconnection requirements, the financial institution contracts its CSP for a high-capacity Carrier Ethernet based DCI service delivered over an optical backbone.

The CSP deployed **RAD's ETX-2i-400G** at all three locations as the demarcation and aggregation platform, creating a seamless, high-performance Layer-2 multipoint DCI fabric interconnecting:

- Both enterprise data centers
- The Azure cloud on ramp Peering Gateway site that extends the private network to Microsoft Azure's GPU training environment

Solution Brief
Data Center Interconnect over Carrier Ethernet

The ETX-2i-400G's integrated **MACsec and quantum safe cryptography** protect sensitive financial and AI training data across all inter sites and Azure bound connections , ensuring compliance with regulatory, sovereignty and audit requirements.

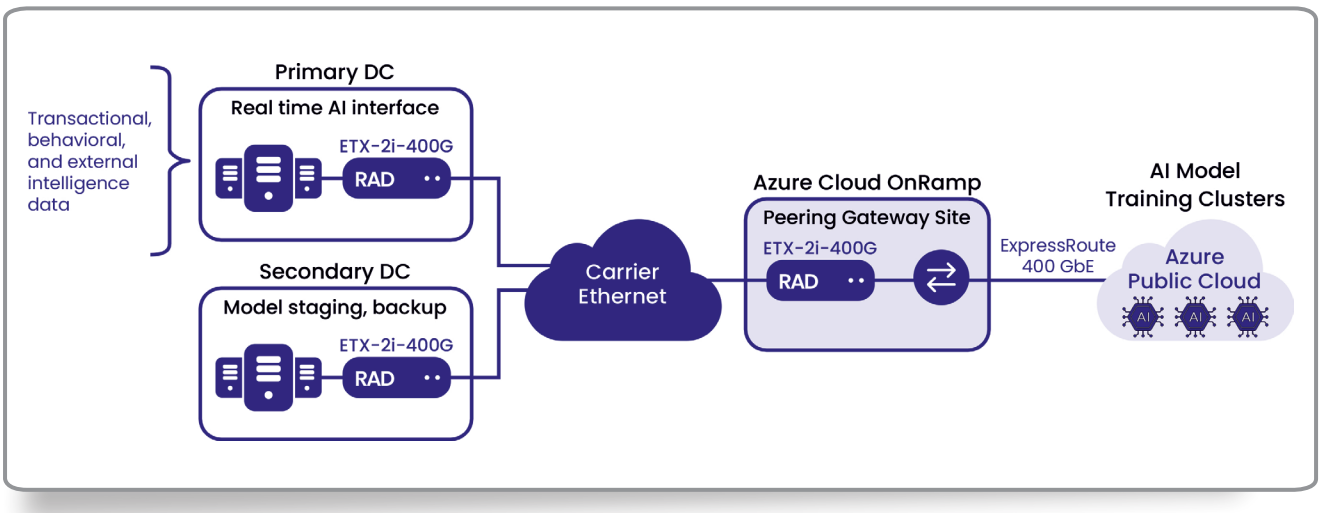

Advanced Traffic Engineering for AI Workloads

With advanced traffic management, HQoS, and per EVC/port shaping, the CSP ensures optimal performance for each workload type:

- **Ultra-low-latency inference flows,** real-time fraud-scoring traffic, and event-driven transaction streams within the DCI fabric are elevated to top-priority queues.
- **High volume Azure bound transfers,** including training datasets, data lake sync, and model checkpointing, are shaped and rate limited to avoid interference with real-time operations.
- **Live transactional data** from branches, ATM or digital channels is correctly marked and prioritized for rapid fraud detection.

The ETX-2i-400G's deep buffers minimize jitter, mitigate microbursts, and prevent congestion, ensuring uninterrupted banking operations and consistent customer experience.

Programmable bandwidth scaling supports ingestion spikes during AI model training cycles, while service assurance provides continuous visibility into latency, jitter, and frame

Solution Brief
Data Center Interconnect over Carrier Ethernet

Reference Use Case Summary

By adopting the CSP's RAD-enabled Layer-2 DCI service, the financial institution benefits from more efficient Azure-based AI training, improved real-time fraud-detection accuracy, and secure, resilient connectivity between its private data centers and its Microsoft Azure environment. At the same time, the CSP expands its portfolio with a scalable service offering designed for hybrid-cloud AI workloads.

For more information on how to improve your Data Center Interconnect offering, contact us at: market@rad.com.

RAD
Your Network's Edge®

Specifications are subject to change without prior notification. The RAD name, logo and logotype, are registered trademarks of RAD Data Communications Ltd. RAD product names are trademarks of RAD Data Communications Ltd. ©2026 RAD Data Communications. All rights reserved. www.rad.com