


CASE STUDY

Intel® Ethernet Converged Network Adapters
10GBASE-T Fibre Channel over Ethernet
City of Pforzheim, Germany



The Next Generation of Converged Networking: Fibre Channel over Ethernet Using 10GBASE-T



The City of Pforzheim, Germany, has a forward-looking IT division that has used virtualization to simplify its data center while meeting the growing demands of city operations. Building on that progress, the organization is now looking to better support its virtualized infrastructure by combining the cost-effectiveness and simplicity of converged LAN and storage traffic with the flexibility of 10GBASE-T.

To modernize its data center, Pforzheim has used VMware vSphere* and Cisco hardware such as Nexus* 5000 Series switches, achieving outstanding results. These virtualization efforts have provided server consolidation that allows for growing application demands to be supported cost-effectively. The city's IT organization has also simplified its environment and realized cost savings by unifying its LAN and storage fabrics using Fibre Channel over Ethernet (FCoE). This approach provides dependable, high-performance access to storage, providing robust support for virtualized usage models based on live migration of virtual machines between physical hosts.

After meeting with Intel engineers at Cisco Live! in 2012, Pforzheim IT decided to test 10GBASE-T with FCoE, using Cisco Nexus 2000 Series fabric extenders and Intel® Ethernet Converged Network Adapters X540-T2. This combination of technologies and products from Intel and Cisco for FCoE over 10GBASE-T proved itself in the Pforzheim IT lab, delivering outstanding performance at lower costs than a comparable Fibre Channel installation. The City of Pforzheim plans to move forward with this topology in its production environment.

CHALLENGES

- **Drive down data center costs and complexity**, even as the demands placed on IT continue to grow.
- **Set the stage for a smooth, gradual transition** from existing solutions while positioning the environment for continuing future growth.

SOLUTIONS

- **Intel® Ethernet Converged Network Adapters and Cisco Nexus switches** unify LAN and storage traffic on a single 10 Gigabit Ethernet fabric.
- **Fibre Channel over Ethernet (FCoE) using 10GBASE-T** represents a novel combination made possible by products and technologies from Intel and Cisco.

RESULTS

- **Met or exceeded Fibre Channel performance at low cost**, without dedicated Fibre Channel switches and adapters.
- **Delivered flexible solution and upgrade path**, including cabling backwards-compatibility, native FCoE initiator support, and a rich industry ecosystem.

Modern Infrastructure for an Ancient City

Pforzheim is a city of about 120,000 citizens near Germany's Black Forest, tracing its roots back some 2,000 years to its origin as a Roman settlement. In modern times, the city is known internationally as a center for jewelry and watch making. The city runs a data center to provide hundreds of applications that support functions in areas such as government, utilities, and schools, as well as services such as telephony. Users of the city's network are distributed among more than 100 buildings, which are connected by a city-owned copper and fiber cabling infrastructure.

Like most IT organizations in both the public and private sectors, Pforzheim must support escalating requirements with a relatively static level of resources. While the increasing costs of personnel, energy, and other needs are outside their control, the people charged with responsibility for the Pforzheim data center strive to build in efficiencies at every turn. Reducing infrastructure requirements using virtualization has been a major contributor to this effort, as has decreasing the challenges associated with managing the environment.

COMPLEMENTARY TECHNOLOGIES FOR NETWORK MODERNIZATION

The need to support ever-increasing IT requirements demands innovation and efficiency. The City of Pforzheim meets those demands using a succession of technology advances, supported by equipment and best practices from providers such as Cisco and Intel:

- **Virtualization** enhanced by server and networking optimizations consolidates servers onto fewer physical hosts while accelerating provisioning, enhancing agility, and enabling flexible usage models.
- **Unified networking** converges LAN and storage traffic onto a single 10 Gigabit Ethernet (10GbE) fabric using Fibre Channel over Ethernet (FCoE) to reduce complexity as well as lower both capital and operating expenses.
- **10BASE-T** provides standards-based, high-bandwidth connectivity that is backwards-compatible with Gigabit networking infrastructures over existing low-cost copper cabling with a reach of up to 100 meters for LAN traffic and up to 30 meters for FCoE traffic.¹

Virtualization has provided significant cost benefits to the city's operations, both by decreasing the number of servers and related equipment needed and by increasing the overall ability to respond to changing needs. In conjunction with virtualizing its data center, converging storage and LAN traffic using FCoE enables virtualized servers to take advantage of the storage area network (SAN) for live migration and related usage models with a single network fabric. This approach delivers the following core benefits:

- **Eliminates the need for dedicated Fibre Channel switches and adapters**, lowering capital costs and energy expense, while enabling standardization on a single adapter for both LAN and storage traffic.²
- **Reduces cabling requirements**, including fewer cables and lower complexity, which decreases expenses and the likelihood of costly human error.
- **Decreases management complexity** because of the change from two fabrics to one, allowing common tools and expertise to support the network.

Ethernet enhancements such as Data Center Bridging (DCB) guarantee successful data delivery by providing for lossless Ethernet, which is instrumental to FCoE. The smooth interoperability between Intel Ethernet Converged Network Adapters and the Cisco Nexus family of switches and fabric extenders is enabled through co-engineering and extensive testing by Cisco and Intel.

While a detailed discussion in this area is outside the scope of this paper, Pforzheim's IT organization also takes advantage of advanced features of the network hardware to enhance redundancy. For example, server ports are teamed using Link Aggregation Control Protocol (LACP), and multipathing is enabled using virtual port channels on the Cisco Nexus switching infrastructure. To optimize performance, the group enabled jumbo frames on the Intel Ethernet Converged Network Adapters, as well as implementing I/O virtualization using virtual machine device queues (VMDq).

Having used FCoE to help meet fiscal and technical challenges, the City of Pforzheim continues to evaluate additional techniques as the basis of innovation that can improve on those gains even further. Pforzheim's IT organization was particularly enthusiastic about taking advantage of the 10GBASE-T support for FCoE from Cisco and Intel. Because passing FCoE traffic over 10GBASE-T connections is a relatively new approach, the city worked with Intel to undertake lab-based testing, which yielded very successful results.

10GBASE-T FOUNDATIONS OPTIMIZE THE BENEFITS OF UNIFIED NETWORKING

With the Intel® Ethernet Controller X540, 10GBASE-T connectivity is now integrated onto mainstream servers through fixed and modular LAN on Motherboard (LOM) designs as well as converged network adapters. The fixed LOM approach offers the benefits of 10 Gigabit Ethernet (10GbE) without additional expense beyond the cost of the core server platform, while modular LOM and converged network adapters based on PCI Express* enhance solution flexibility. All 10GBASE-T implementations offer the following benefits:

- **Backwards-compatibility with Gigabit Ethernet (GbE)**, providing auto-negotiation between GbE and 10GbE connections that makes the transition to 10GbE relatively seamless and allows for phased migration.
- **IEEE standard-based compatibility with existing cabling**, including cost-effective Category 6A (CAT6A) twisted-pair copper (the minimum requirement for Fibre Channel over Ethernet (FCoE) traffic), as opposed to the more costly, non-standard, 10GbE-only cabling infrastructure when using SFP+ Direct Attach.
- **Extended cable reach for the flexibility to support multiple deployment models**, at distances of up to 100 meters for LAN traffic and up to 30 meters for FCoE traffic.¹



Test Environment for Unified Networking over 10GBASE-T

The environment that the City of Pforzheim used to verify the viability of 10GBASE-T networking for FCoE traffic includes the building blocks shown in Table 1. The servers connect to the Cisco Nexus 2232TM-E Fabric Extenders using 10GBASE T, which forward traffic to the Cisco Nexus 5548P Switch over 10GbE SFP+ uplinks. The fabric extenders function as remote line cards for the parent switch, forming a distributed modular system where traffic is shaped according to policies established on the parent switch.

Table 1. Test environment for passing FCoE traffic over 10GBASE-T.

Servers	<p>Approximately 80 physical hosts:</p> <ul style="list-style-type: none"> ▪ Intel® Server Systems R2208GZ4GC: two Intel® Xeon® processors E5-2690, 258GB DDR3-1600 reg. ECC Memory, RMS25PB080 RAID-Module ▪ HP ProLiant* DL380 G7 Server: two Intel® Xeon® processors X5690, 192GB DDR3 reg. ECC Memory, Smart Array Raid-Module
Network Adapters	<ul style="list-style-type: none"> ▪ Intel® Ethernet Converged Network Adapter X540-T2 (dual-port 10GBASE-T adapter) ▪ Dual RJ-45 port 10GBASE-T IO Module AXX10GBTW/LHW based on the Intel® Ethernet Controller X540
Fabric Extenders	Cisco Nexus* 2232TM-E
Network Switch	Cisco Nexus 5548P
Storage Targets	<ul style="list-style-type: none"> ▪ One NetApp FAS3140* storage system ▪ Two NetApp FAS3240* storage systems
Virtualization Environment	<p>VMware vSphere* 5.1 (Guest OS: Microsoft Windows Server* 2008 R2)</p>

The distributed modular system architecture provides a single point of management for the entire server infrastructure, and the relatively low cost of the fabric extenders enhance the cost-effectiveness of the infrastructure as a whole. The fabric extenders also provide high density on the rack, with 32 host interfaces and 10 uplink ports per 1RU module, and expandability with capacity for up to 24 fabric extenders per switch.

Successful Proof of Concept with FCoE and 10GBASE-T

Performance testing with FCoE traffic and 10GBASE-T connectivity revealed that this architecture easily meets the City of Pforzheim's requirements. To measure and characterize I/O performance of the test environment, the team used two test applications:

- **lometer**, an I/O subsystem measurement and characterization tool for single and clustered systems that was originally developed by Intel and that is now community-maintained as an open-source project.
- **fio (flexible I/O tester)**, an open-source toolkit to simulate I/O loads for benchmarking or stress testing by spawning a configurable number of threads or processes, each doing a particular type of I/O action.

Testing was conducted across a range of packet sizes, including 4 KB, 8 KB, 16 KB, and 64 KB; specifically, smaller packet sizes were used to measure maximum I/O per second, and larger packets were used to measure I/O bandwidth. In production environments, typical implementations might use 4-KB packets for transaction-based systems, while 64-KB packets might be used for file and print systems. The testing used fio to spawn multiple processes that simulated I/O loads, such as simulated Oracle Database* loads with various block sizes and direct (non-buffered O_DIRECT) and libaio-based asynchronous I/O. Performance results as measured by lometer are shown in Figure 1.




Figure 1. Performance results for Fibre Channel over Ethernet using 10GBASE-T, obtained using lometer.

lometer testing achieved 100,508 I/O operations per second at 4-KB packet sizes, and throughput of approximately 1.05 GB per second with 64-KB packets. In addition to meeting the performance requirements of Pforzheim IT, the 10GBASE-T topology maintained an error rate of zero, in keeping with the needs and expectations of the Pforzheim IT organization for lossless transmission of storage traffic. Tests conducted using fio achieved similar results.

“Our testing with Intel® Converged Network Adapters and Cisco Nexus* switches provided excellent performance without any errors using 10GBASE-T for Fibre Channel over Ethernet.”

– Andreas Hurst, Head of Information and Communications Technology, City of Pforzheim

These results provide the substrate for live migration and related virtualized usage models that depend on robust, dependable SAN access. In particular, this testing demonstrated to the City of Pforzheim team that FCoE using 10GBASE T is a viable alternative for many of the workloads presently supported by Fibre Channel in the city's data center environment. Those workloads include storage traffic for Windows Server, Microsoft Exchange*, Oracle Database, Microsoft SQL Server*, and data backup to the city's Linear Tape-Open (LTO) magnetic tape storage library.



Conclusion

Based on the success of its testing, the City of Pforzheim has decided that 10GBASE-T is a viable solution for deployment in their data center, including as a networking technology for FCoE. "Our testing with Intel Converged Network Adapters and Cisco Nexus switches provided excellent performance without any errors using 10GBASE-T for Fibre Channel over Ethernet," said Andreas Hurst, Head of Information and Communications Technology for the City of Pforzheim. "This approach is an excellent, more cost-effective alternative to native Fibre Channel, and 10GBASE-T promises to make broad deployment of 10 Gigabit Ethernet in our data center faster and easier. Intel and Cisco have also given our team excellent support during this process."


The city is now planning a widespread deployment of 10GBASE-T, replacing some existing Fibre Channel infrastructure and providing a cost-effective, phased migration to 10GbE. This implementation is expected to take advantage of existing cabling infrastructure, policies, and expertise while enhancing the municipality's ability to deliver services that meet or exceed the growing demands of its citizens.

To learn more about Intel® Ethernet Converged Network Adapters, visit:
www.intel.com/go/ethernet

To learn more about Cisco Nexus Switches and Fabric Extenders, visit:
www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html

SOLUTION PROVIDED BY:

IN COLLABORATION WITH:



¹ Tests conducted in June 2013 by Intel demonstrated a bit error rate of better than 10^{-15} over 30 meter, 2-connector channels using Category 6A unshielded twisted-pair cabling, Category 6A foiled unshielded twisted-pair cabling, and Category 7 foiled unshielded twisted-pair cabling hardware, cabling and patch cords.

Test configurations:

Intel® Xeon® Processor 5150, 2.66 GHz, 2GB-3.25GB RAM, Intel Ethernet Converged Network Adapter X540-T2. Supermicro SuperServer* 7046T-3R server with Supermicro X8DT3 motherboard and Intel® (Tylersburg) chipset, Cisco Nexus 3064-T 13x port-adapter pairs, 500M (minimum) 1518 byte packets transmitted from adapter to switch, 100x trials/pair, Panduit PUR6A04 TX6A CMR (UTP), PUF6X04 TX6A CMR (U/FTP) and PSL7004 LSZH (S/FTP) 30m channels.

Intel® Xeon® Processor 5440, 2.83 GHz, 2GB-8GB RAM, Intel Ethernet Converged Network Adapter X540-T2. Supermicro SuperServer* 7045W-NTR+B server with Supermicro X7DWN+ motherboard and Intel® 5400 (Seaburg) chipset, Cisco Nexus 7000 F2-Series N7K-F248XT-25 13x port-adapter pairs, 500M (minimum) 1518 byte packets transmitted from adapter to switch, 100x trials/pair, Panduit PUR6A04 TX6A CMR (UTP), PUF6X04 TX6A CMR (U/FTP) and PSL7004 LSZH (S/FTP) 30m channels.

Tests conducted using Intel's NEXBITT interoperability test software and Intel's NPPV remote networking controller hardware test tool. All switch and adapter endpoints were configured to autonegotiate link speed and duplex mode.

² Based on using a single FCoE-capable switch for LAN and SAN traffic.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

Results have been measured by Intel based on benchmark or other data of third parties and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. Intel does not control or audit the design or implementation of third party data referenced in this document. Intel encourages all of its customers to visit the websites of the referenced third parties or other sources to confirm whether the referenced data is accurate and reflects performance of systems available for purchase.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families: Learn About Intel® Processor Numbers:

www.intel.com/content/www/us/en/processors/processor-numbers.html.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice revision #20110804

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/performance.

*Other names and brands may be claimed as the property of others.

Copyright © 2013 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

0713/BY/MESH/PDF

328793-001US