



NAS or iSCSI?

Selecting a storage system

White Paper 2006



Contents

“iSCSI or NAS: selecting a storage system”

1.	A growing market: Network storage 4
2.	Introduction to storage systems 4
2.1	Storage concepts 5
2.1.1	Direct Attached Storage 5
2.1.2	Network Attached Storage 5
2.1.3	Storage Area Network 5
3.	Decision support for choosing a storage system: NAS or iSCSI? 6
3.1	Typical NAS applications (Network Attached Storage) 6
3.1.1	Advantages 6
3.1.2	Disadvantages 7
3.2	Typical iSCSI applications (Internet Small Computer Systems Interface) 8
3.2.1	Advantages 9
3.2.2	Disadvantages 10
3.3	Summary 10
4.	Open-E: Storage software for any requirement 11
4.1	Open-E NAS product family 11
4.2	Open-E iSCSI product family 11
5.	About Open-E: Guaranteed storage economy 11

©2006 by Open-E GmbH. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical, photocopying or otherwise, without the prior written agreement or permission of Open-E GmbH, Lindberghstr. 5, D-82178; Puchheim, Germany.

Open-E is a registered trademark of Open-E GmbH. All other trademarks mentioned herein are property of their respective owners.

Open-E assumes no responsibility for errors or omissions in this document, nor does Open-E make any commitment to update the information contained here in. Any information contained herein may change at any time without prior notice.

A growing market: Network storage

- Worldwide, over 50 petabytes (5×10^{16} bytes) of data are generated every day that have to be safely stored, economically managed, and quickly and efficiently made available to applications and users when required. In 2004, this steadily expanding data volume amounted to 1.5 exabytes (1.5×10^{18} bytes) - and it is still growing.
- One storage heavyweight, for example, is the San Diego Supercomputer Center (SDSC) in the USA, which stores approximately one petabyte of data on hard drives and archives another six petabytes on tape. Another example is the Deutsches Elektronensynchrotron (DESY) in Hamburg (Germany), which generates 100 terabytes of raw data in each experiment. But also in everyday work, hundreds of terabytes of storage volume are generated: The data centre of a local hospital like the Klinikum Oldenburg in Germany, has to deal with 20 gigabytes of data every month in form of electronic patient files.
- So it's no surprise that market forecasters are predicting high growth rates for storage solutions. Figures from IDC's latest market research show: From 2005 to 2008, sales (in USD) of **NAS** will increase annually by almost 15%, and sales of **iSCSI** solutions annually by an impressive 170%. At an increase of 785% over the entire period, the data volume for **NAS**- and **iSCSI**-based low-cost disk systems will increase almost eight-fold!
- IDC expects the market share of **NAS/iSCSI** to grow from 13.7% of sales in 2005 to 30.2% of sales in the year 2008. In terabytes, **NAS/iSCSI**'s market share will increase even from 24.9% to 42.5%!
- Thanks to more powerful drives with higher capacities, the price per Gbyte of storage will decrease by 30% every year. So while sales in the low-cost disk systems market (**NAS/iSCSI**) will be growing by 190% annually between 2005 and 2008, the data volume on **NAS** will be growing by 370%, and that on **iSCSI** by a stunning 2,475%!

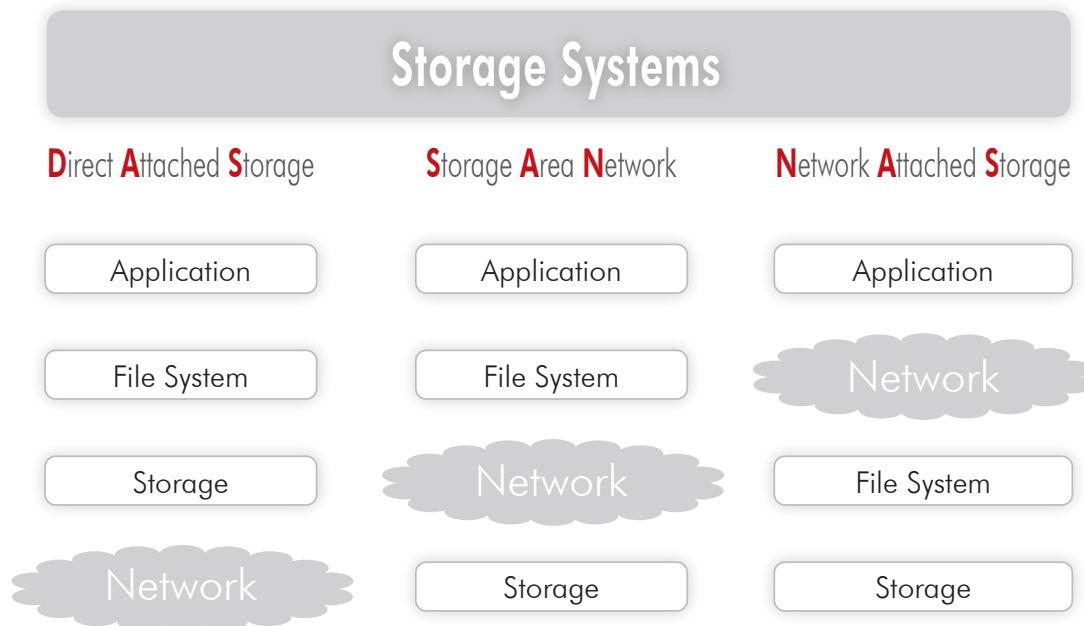
Introduction to storage systems

DAS, **NAS** and SAN differ mainly in where they interface with the network.

- There are currently three basic concepts for addressing storage:
 - Direct Attached Storage (DAS),
 - Network Attached Storage (**NAS**),
 - Storage Area Networks (SAN).
- De facto, no computer system can do without mass storage – mainly hard drives. But how these storage media are connected has changed drastically over the course of their development.



Introduction to storage systems



Direct Attached Storage

- Direct Attached Storage is a system of hard drives addressed directly via system buses within the computer (IDE, SCSI); the network interface is managed by the operating system. As these buses can only bridge short distances within the decimeter range, DAS solutions are limited to the respective computer casing. Depending on the bus type, DAS systems are also restricted to a relatively small number of drives - Wide-SCSI achieves the maximum of 16 directly addressable drives. Due to these limitations and the need for more flexible storage, the importance of DAS is declining. Although DAS in terms of terabyte is still growing by 28% annually, the need for storage is increasingly being covered by networked storage like **NAS** and **iSCSI** systems.

Network Attached Storage

- NAS systems are generally computing-storage devices that can be accessed over a computer network (usually TCP/IP), rather than directly being connected to the computer (via a computer bus such as SCSI). This enables multiple computers to share the same storage space at once, which minimizes overhead by centrally managing hard disks. **NAS** devices become logical file system storage for a local area network. **NAS** was developed to address problems with direct attached storage, which included the effort required to administer and maintain “server farms”, and the lack of scalability, reliability, availability, and performance. They can deliver significant ease of use, provide heterogeneous data sharing and enable organizations to automate and simplify their data management.

Storage Area Networks

- Storage Area Networks (SAN), which also include **iSCSI**, are distinguished from other forms of network storage by using a block-based protocol and generally run over an independent, specialized storage network. Data traffic on these networks is very similar to those used for internal disk drives, like ATA and SCSI. With the exception of SAN file systems and clustered computing, SAN storage is still a one-to-one relationship. That is, each device (or Logical Unit Number (LUN)) on the SAN is “owned” by a single computer (or initiator). SANs tend to increase storage capacity utilization, since multiple servers can share the same growth reserve. Other benefits include the ability to allow servers to boot from the SAN itself. This allows for a quick and easy replacement of faulty servers since the SAN can be reconfigured so that a replacement server can use the LUN of the faulty server.

Selecting a storage system: **NAS** or **iSCSI**?

- This white paper's aim is to help system administrators decide between the two storage technologies **NAS** and **iSCSI**. Both technologies were developed to help cover the exponentially increasing need for storage capacity by means of external, scalable storage devices that can be networked if necessary.
- In any company, the use of storage solutions will depend on a number of different factors:
 - Available budget
 - Data security requirements
 - Network infrastructure
 - Data availability requirements, etc.
- Within these basic considerations, system administrators must decide which technology and which products are best suited to meet their needs.

Typical **NAS** applications

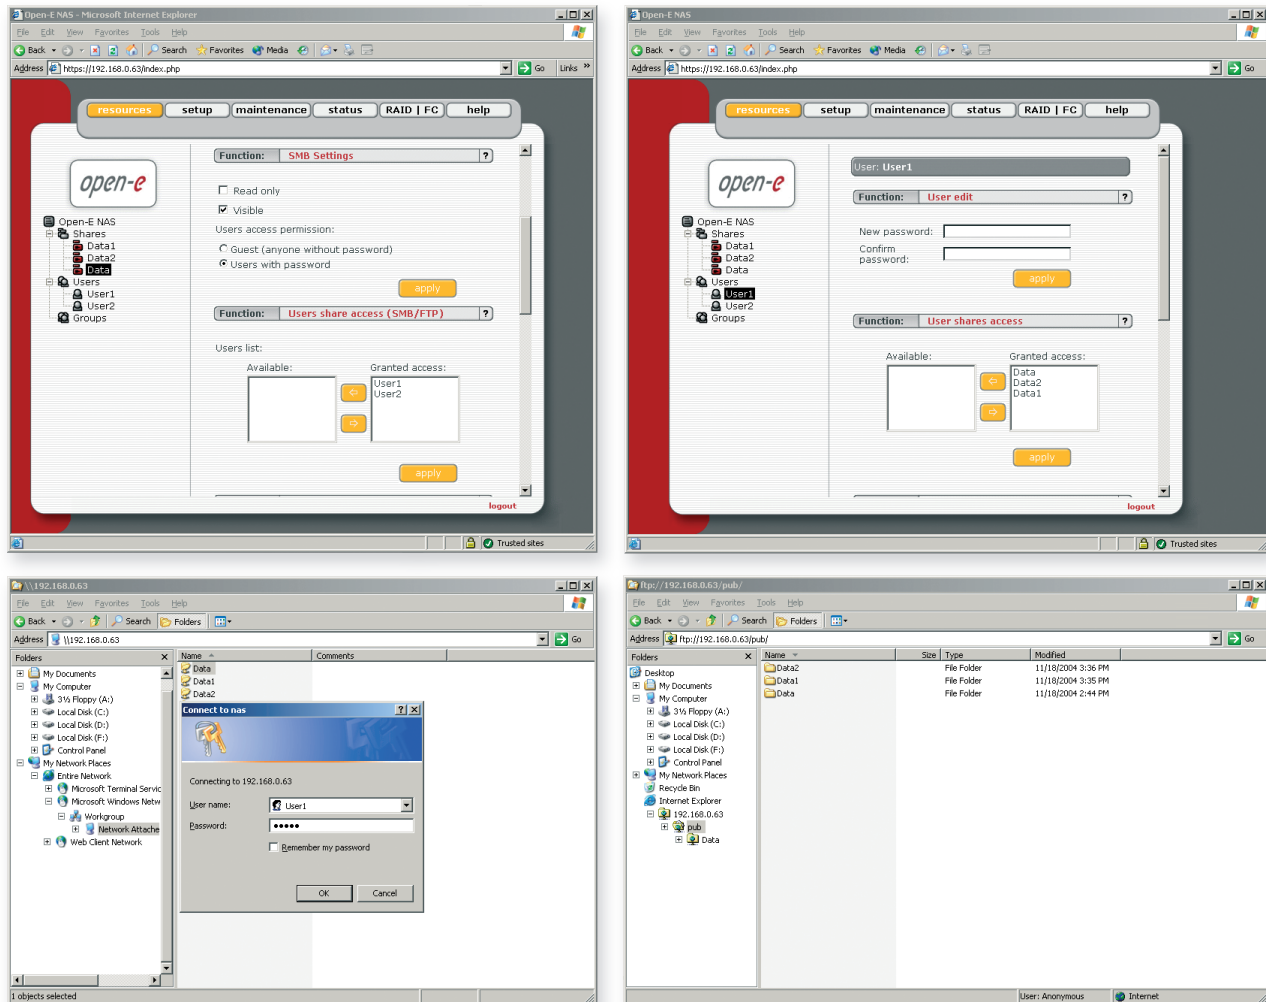
- A **NAS** server is a specialized appliance optimized for file serving, i.e., storing, retrieving, and serving files. Typical applications include:
 - File/Print server
 - Application specific server
 - Video Imaging
 - Graphical image store
 - Centralized heterogeneous file sharing
 - File system mirroring
 - Snap shot critical data
 - Replacement of traditional backup methods
 - Medical imaging
 - CAD/CAM
 - Portable centralized storage for offsite projects
 - Onsite repository for backup data

Advantages of **NAS**

- **NAS** systems offer a number of advantages:
 - Heterogeneous OS support. Users running different types of machines (PC, Apple iMac, etc.) and running different types of operating systems (Windows, Unix, Mac OS, etc.) can share files.
 - Easy to install and manage. **NAS** appliances are "plug-and-play" meaning that very little installation and configuration is required beyond connecting them to the LAN.
 - **NAS** appliances can be administrated remotely, i.e. from other locations.
 - Less administration overhead than that required for a Unix or Windows file server.
 - Leverages existing network architecture since **NAS** are on LANs.
 - **NAS** server OSs are smaller, faster, and optimized for the specialized task of file serving and are therefore undemanding in terms of processing power.
 - A **NAS** appliance is a standalone file server and can free up other servers to run applications. Compared to **iSCSI** an additional host server is not necessary.
 - Compared to **iSCSI**, **NAS** appliances already include integrated mechanisms for backup, data synchronization and data replication.



Typical NAS applications



The screen shots show user authorization settings on the web-based graphic user interface of Open-E NAS and secure access to NAS storage from a Windows client.

What are some disadvantages of NAS?

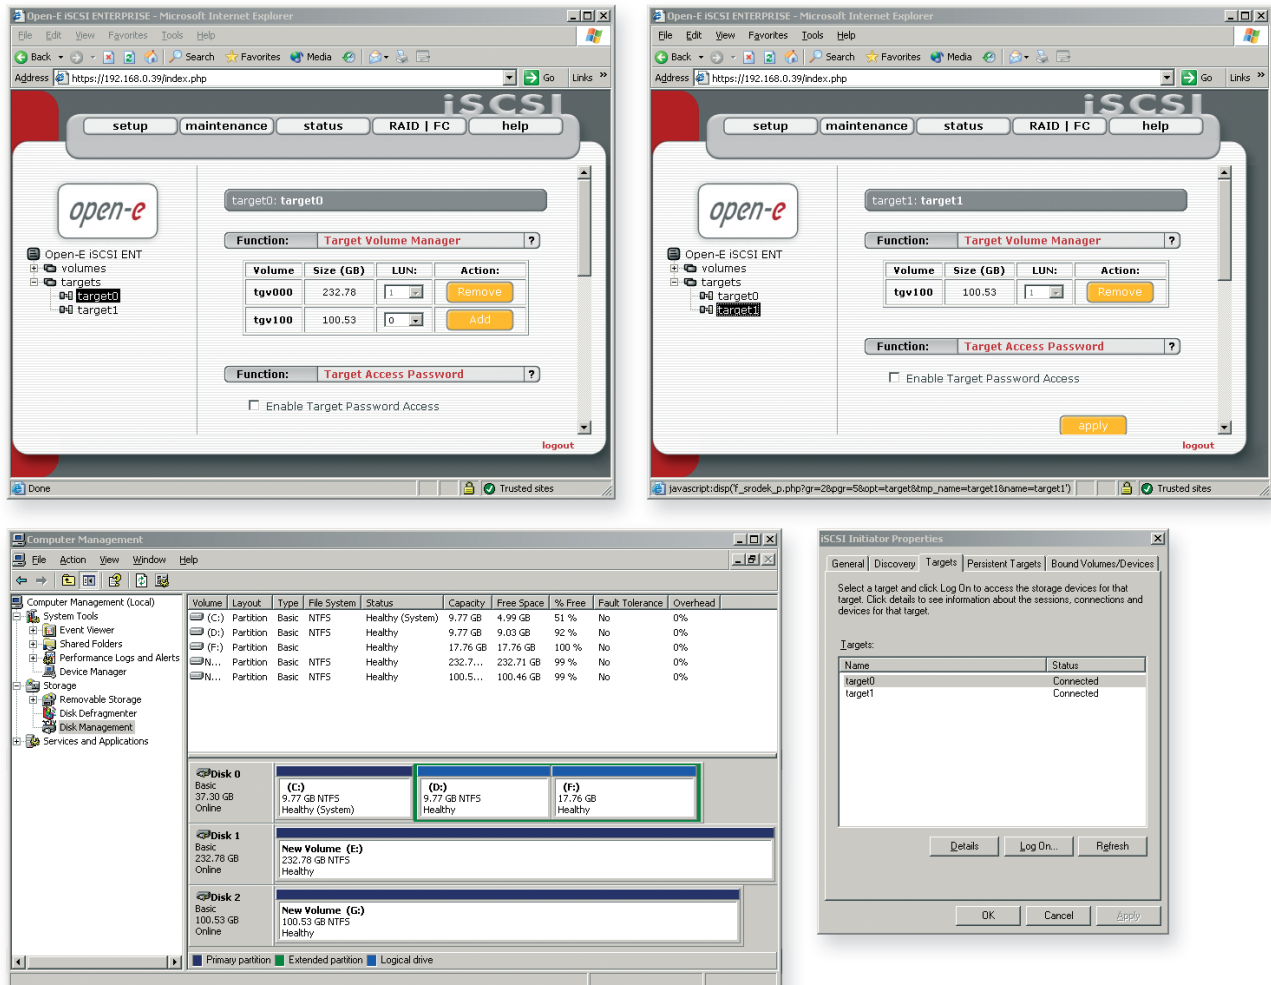
- NAS has the following disadvantages:
 - Heavy use of NAS will clog up the shared LAN negatively affecting the users on the LAN. Therefore NAS is not suitable for data transfer intensive applications.
 - Somewhat inefficient since data transfer rides on top of standard TCP/IP protocol.
 - Cannot offer any storage service guarantees for mission critical operations since NAS operates in a shared environment.
 - NAS is shared storage. As with other shared storage, system administrators must enforce quotas without which a few users may hog all the storage at the expense of other users.
 - NAS is less flexible than a traditional server.
 - Most database systems such as Oracle or Microsoft Exchange are block-based and are therefore incompatible with file-based NAS servers (except for SQL).

Typical iSCSI applications

- iSCSI stands for Internet Small Computer System Interface, a technology that extends the Internet Protocol to Storage Area Networks (SAN).
- A primary usage of iSCSI is instant disk space. With no server downtime or re-boot required you install the initiator driver on f.e. the Windows machine, map the storage to the disk then allocate the desired disk capacity and format it under disk manager.
- An iSCSI SAN can utilize an existing Gigabit network to serve up storage and is therefore particularly appealing for small and medium-sized businesses. It leverages the investment in the networking infrastructure and also the storage management, as most IT professionals understand TCP/IP. Adding Fibre Channel storage involves extensive training on mapping and managing storage.
- Of course NAS systems can also be implemented via iSCSI, but from a financial point of view this only makes sense if an iSCSI-based storage solution is to be set up anyway. On the other hand, iSCSI targets can also extend the storage capacity of a NAS server - although this requires an iSCSI initiator on the NAS server.
- Typical iSCSI applications include:
 - Offer power users disk space on demand
 - Run databases on iSCSI disk
 - Video Imaging
 - Graphical image store
 - File system mirroring
 - Snap shot critical data
 - Replacement of traditional backup methods
 - Medical imaging
 - CAD/ CAM
 - Onsite repository for backup data



Typical iSCSI applications



A striking advantage: Within an operating system - here Windows XP - iSCSI displays storage space on the network as a local drive, and also allows it to be managed like one.

Advantages of iSCSI

- iSCSI systems offer a number of advantages:
 - Ease of scaling disk storage. With iSCSI the disks are remote from the server, therefore adding a new disk just requires the use of disk manager or if replacing the whole server re-mapping the data to the server using iSCSI. With iSCSI you can easily create huge storage pools with volumes in the range of several tera- or petabytes.
 - In comparison to NAS, which provides a file-level interface, iSCSI provides a block level interface and are therefore compatible with database applications such as Oracle or Microsoft Exchange, that also use a block level interface.
 - Leverages existing network architecture since iSCSI are on LANs.

Typical iSCSI applications

Advantages of iSCSI

- iSCSI storage appliances can be seamlessly integrated into existing SAN environments, since it also runs on block level storage.
- iSCSI can provide significant benefits for providing failover for high availability configurations. iSCSI allows IP based replication, mirroring and clustering of data and offers integrated MPIO (Multi-Path I/O).
- iSCSI can also be configured as a particularly flexible DAS system – the local SCSI bus is so to speak extended by the network. As iSCSI
- is an underlying technology to the OS and uses the native file system of the applications, it is fully compatible with all file systems.

What are some disadvantages of iSCSI?

- The demands of accommodating SCSI commands and SCSI data packets in TCP/IP packets require extensive hardware resources: CPU performance should be at least that of a 3 GHz Pentium processor, Gigabit Ethernet (GbE) should accordingly be used as a network interface, and the RAM requirement is also significant.
- For sharing iSCSI targets with multiple initiators, additional server or specific (client) software for shared data access is required. Known providers of SAN data sharing software are Adic, Sanbolc, IBM, Polyserve, Dataplow and SGI.
- As iSCSI is an underlying technology to the OS and application, anything that an organization currently has can be used. On the other hand this means that extra licenses for OS and software applications might be needed.
- Comparing to NAS, an iSCSI Target is not a standalone device and an additional host server is necessary.
- For sharing centralized storage pool disk among heterogeneous OS requires additional sharing software.
- In iSCSI appliances mechanisms for f.e. backup, data synchronization and data replication are not integrated and must be configured. Comparing to NAS, iSCSI behaves like a local hard drive in the network.

Summary

- Network Attached Storage (NAS) is the obvious choice for a storage solution wherever the main focus is on storing and archiving files and shared access to these over a network – even from different client operating systems. Small and medium-sized businesses, typing pools, legal or agency offices, and even end users with large amounts of multimedia files will find an affordable storage solution for their needs in NAS.
- For storing database systems - other than SQL-based database systems - on a network, NAS is however not a feasible solution. For requirements of this type the industry has developed the Storage Area Network (SAN) technology, which can often be implemented using iSCSI components. Advantages of iSCSI: An IP-based SAN allows administrators to use their familiar management tools and security mechanisms and rely on their existing know-how. However, iSCSI only makes sense in connection with a fast LAN infrastructure: At a throughput of approximately 120 Mbyte/s, the performance of a 1 Gbit Ethernet will be sufficient for database applications for approximately 100 users (data volume: approx. 15 MByte/s). Only high-end storage systems will require a 10 GbE infrastructure. Somewhat inefficient since data transfer rides on top of standard TCP/IP protocol. In contrast, SAN uses protocols designed especially for data transfer (though the advantage disappears if a server on the LAN is used to provide a file interface to a SAN).



Open-E: Storage software for any requirement

- With its Open-E **NAS** and Open-E **iSCSI** product lines, Open-E offers system administrators affordable storage solutions that allow them to cover almost any contemporary storage system requirement.

Open-E **NAS** product family

- The Open-E **NAS** family consists of the Flash-based software modules SOHO, 2.0 and Enterprise, which can all be attached to the primary IDE port of a computer in place of a system hard drive.
- The SOHO version is suitable for small installations with no hardware RAID or backup software support to directly address hard drives via IDE or SATA. Version 2.0 covers typical file sharing scenarios in environments with medium failure tolerance. It already offers snapshot functionality and supports all the main hardware RAID controllers. At the high end is the ENTERPRISE version for high availability scenarios with one of the most complete feature sets in the **NAS** segment, including load balancing, UPS support, support for all common backup systems, multi-layered access control and disk quota control.

Open-E **iSCSI** product family

- The Open-E **iSCSI** product family consists of the Flash-based software modules SOHO, SMB and Enterprise, which can all be attached to the primary IDE port of a computer in place of a system hard drive.
- The SOHO version is one of the most affordable ways to quickly and easily deploy an **iSCSI** device. The SMB version is suitable for medium requirements regarding high availability and functionality, providing snapshot function, support for hardware RAID controllers and 10 Gigabit Ethernet NICs, including those with offload hardware. The ENTERPRISE high-end version on the other hand has been optimised for highest performance in high-security environments. It can handle several network connections, adaptive load-balancing and several access control layers, and offers IPSec-, multi-processor and UPS support. For use in professional enterprise SANs the Enterprise version of Open-E **iSCSI** even offers fibre channel support.

About Open-E

Open-E – Guaranteed storage economy

- Open-E was founded on 9th September, 1998, in Bremen, Germany. In late 2001 a dedicated software team for developing storage products was formed at Open-E. Our many years of professional experience in supplying, supporting and providing technical consulting for the use of storage products showed us that there was a gap in the market for tools enabling small, medium and large companies to meet the growing demand for storage solutions. Open-E GmbH in Puchheim near Munich is a software company that develops storage software enabling system integrators to set up high-performance, easy-to-use storage systems. No special knowledge is necessary to install them: Open-E software is compatible with all common operating systems and **iSCSI** initiators, so tailor-made storage systems for companies of all sizes can be implemented with our two product lines Open-E **NAS** and Open-E **iSCSI**.



The information contained in this document represents the current view of Open-E GmbH on the issues discussed as of the date of publication. Because Open-E must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Open-E, and Open-E cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. OPEN-E MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Open-E GmbH.

Open-E may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Open-E, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2006 Open-E GmbH. All rights reserved.

Open-E, Open-E NAS are either registered trademarks or trademarks of Open-E GmbH in Germany and/or other countries.