Harald Seipp
Consulting IT Specialist, Center of Competence for OpenStack Storage Leader
IBM EMEA Storage Competence Center

IBM Software Defined Storage and ownCloud Enterprise Edition

A perfect match for Enterprise File Sync&Share

Based on material by Oliver Oberst, Axel Köster, ownCloud Inc. and pro-com

Project Overview

- Requirements:
  - Efficient and secure storage and application server infrastructure for large scale ownCloud installations

- Proof-of-Concept (PoC):
  - Together with ownCloud Inc. and Business Partner pro-com in Sep 2013 due to a specific customer request

  - Goals:
    - Prove scalability by using technology offering code-named Elastic Storage based on GPFS technology as file system
    - Gather sizing information for six digit user number scenarios
  - Method:
    - Estimate average usage scenario
    - Design architecture
    - Measure performance according to usage scenarios
    - Extrapolate measurement

Blueprint

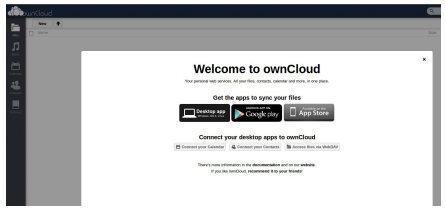

Clients

- Web:
  - Browser
- Desktop Sync:
  - Win
  - Mac
  - Linux
- Mobile:
  - Android
  - iOS

https
webDav

IBM Flex System (Managability)
or IBM NeXtScale (Cost-optimized)
80% OwnCloud Application Servers
20% OwnCloud Database Servers

GPFS

Infiniband

GPFS Storage Server (GSS)
or Elastic Storage Server
OwnCloud User Data

		Model 24	Model 26
		4 JBODs	6 JBODs
	TB	x2/3/4 TB	x2/3/4TB
x24	x48	x232	x348

PoC Setup – Software

- System Software:
  - OS on all Servers: RedHat Enterprise Linux (RHEL) 6.4
  - Web Server:
    - Apache 2.2.15
    - PHP 5.3.3
  - Database:
    - MySQL 5.1.66
  - ownCloud Enterprise Edition 5.0
  - Active ownCloud apps:
    - Deleted files, First Run Wizard, Image Viewer, Provisioning API, Share Files, Text Editor, User Account Migration, Version, ownCloud Instance Migration

PoC Setup – Benchmarks

- **oc-stress:**
  - PHP script which uses curl to generate parallel load on an ownCloud instance
  - Example: `./oc-stress.php GET url 100000 200`
    - 200 concurrent curl requests
    - 100000 requests in total

- **ab**
  - Apache benchmark tool
    - can be used for measuring round-trip speed of single requests if used without concurrency mode
  - Example: `./ab -n100000 -c200 url`
    - 200 concurrent curl requests
    - 100000 requests in total

PoC Procedure

- PoC Procedure:
  - Infrastructure setup
  - ownCloud Enterprise Edition installation
  - Basic/standard performance tuning:
    - e.g. Apache, MySQL
  - First measurement iteration
  - Fine tuning:
    - Mainly OS (IPv4/TCP)
  - Second measurement iteration
  - Optional tuning
  - Final measurement of the major test scenarios

- Major test scenarios due to typical ownCloud workflow scenarios:
  - HTTP_PROPFIND requests (sync client heartbeat)
  - List of directories via WebDAV
  - File upload
  - File download

PoC Measurement Results I

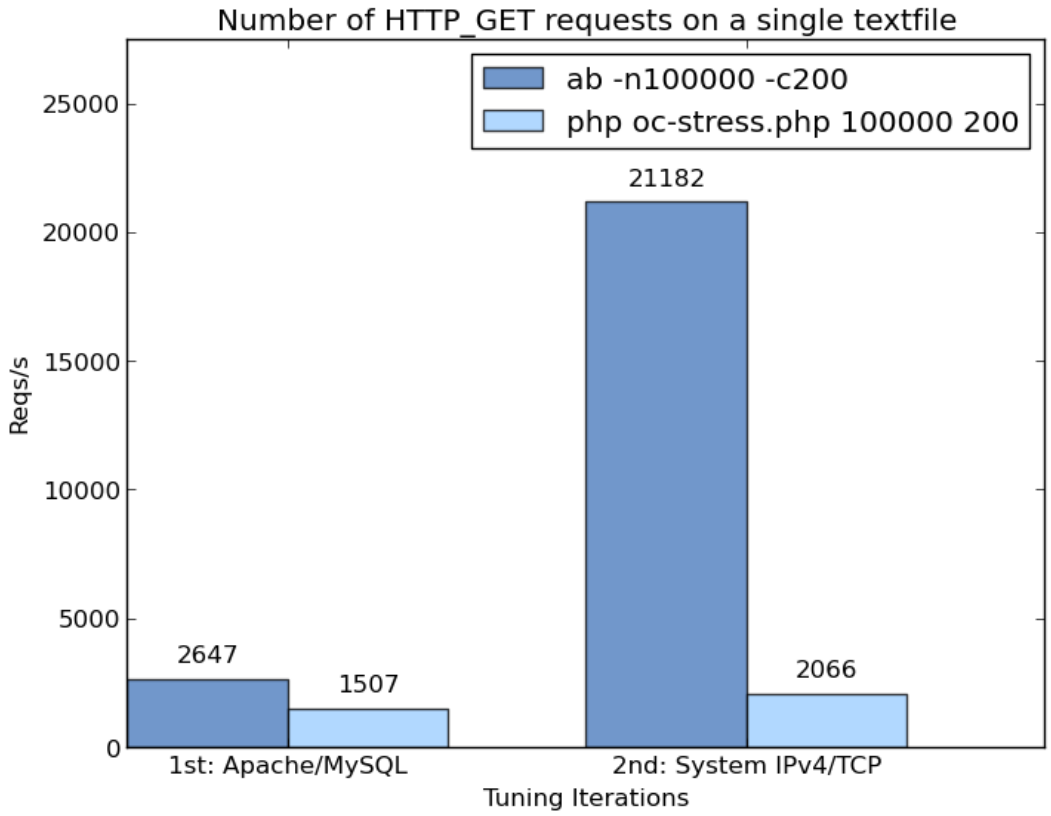

- Results of the Benchmarks after the first and second tuning iteration.
- Single Server/Client
- Results do not converge after tuning the network system parameters. ab just scales better than oc-stress on the benchmark client
- Both show significant improvements.

PoC Measurement Results II

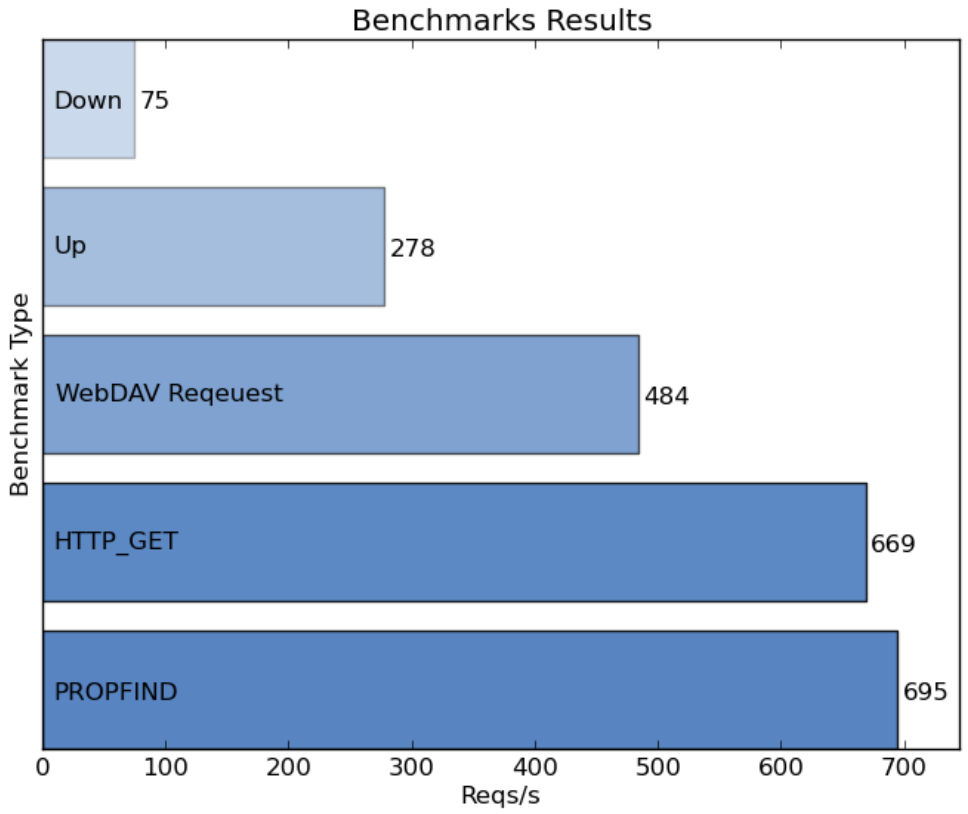

- Final results of the different application server functions of interest within an ownCloud instance
- Each function was executed separately with a concurrency of 1000 parallel requests except for up and download where 100 parallel requests were committed
- Single application server view
- Server load-balance:
  - 80%(App) / 20%(DB)
- Load on GSS at ~5% fluctuation
  - At Idle state as well as during ownCloud benchmark
    → No real impact

PoC Measurement Interpretation

- Basic results:
  - 90% of the load in a typical usage scenario is PROP_FIND (Desktop-Syncing)
  - GSS: only very little additional load on the servers could be seen throughout the measurements
  - Bottleneck: application and database servers
- Projected on a 100.000 user scenario:
  - One desktop client creates 1 PROPFIND / 30sec (Interval configurable)
    - 3333 PROPFIND / sec
      → matched by 5 Nodes
  - One mobile sync per user
    - 28 PROPFIND / sec + 28 GET / sec
      → matched by 0.25 Node
  - One upload and two download per user per hour:
    - 56 downloads / sec + 28 uploads / sec
      → matched by 1.5 Nodes
- Overall estimation for 100k users:
  - 6.75 application servers + 1.35 database servers = 8.1 nodes + 20% buffer
    → approx **10 nodes/100k users** (+ 1 GSS 24/26)

Details: TechDocs Whitepaper

ownCloud Enterprise Edition on IBM Infrastructure

A Performance and Sizing Study for Large User Number Scenarios

Dr. Oliver Oberst – IBM
Frank Karlitschek – ownCloud

Version 5/22/2014
http://w3.ibm.com/support/Techdocs
ownCloud Enterprise on IBM Infrastructure - A Performance and Sizing Study
Page 1 of 11

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102453

https://owncloud.com/wp-content/uploads/ownCloud-on-IBM-Infrastructure.pdf

Solution Architecture, a Client Example

- ownCloud solution for 3 sites
- Storage Backend GSS26
  - 3/1/1 at the 3 sites w/ 1PB net capacity each, 8+3P redundancy
- Frontend Server
  - Per Site:
    - 16 x Application Node nx360 M4
    - 4 x Database Node nx360 M4
    - 2 x Ethernet Load Balancer x3550 M4
- 1 x Management Server nx360 M4
- 2 19“ Racks per site
- Networking components (IB, 10GigE)
- Software
  - ownCloud Enterprise, up to 300k users

Architecture – Data Flow

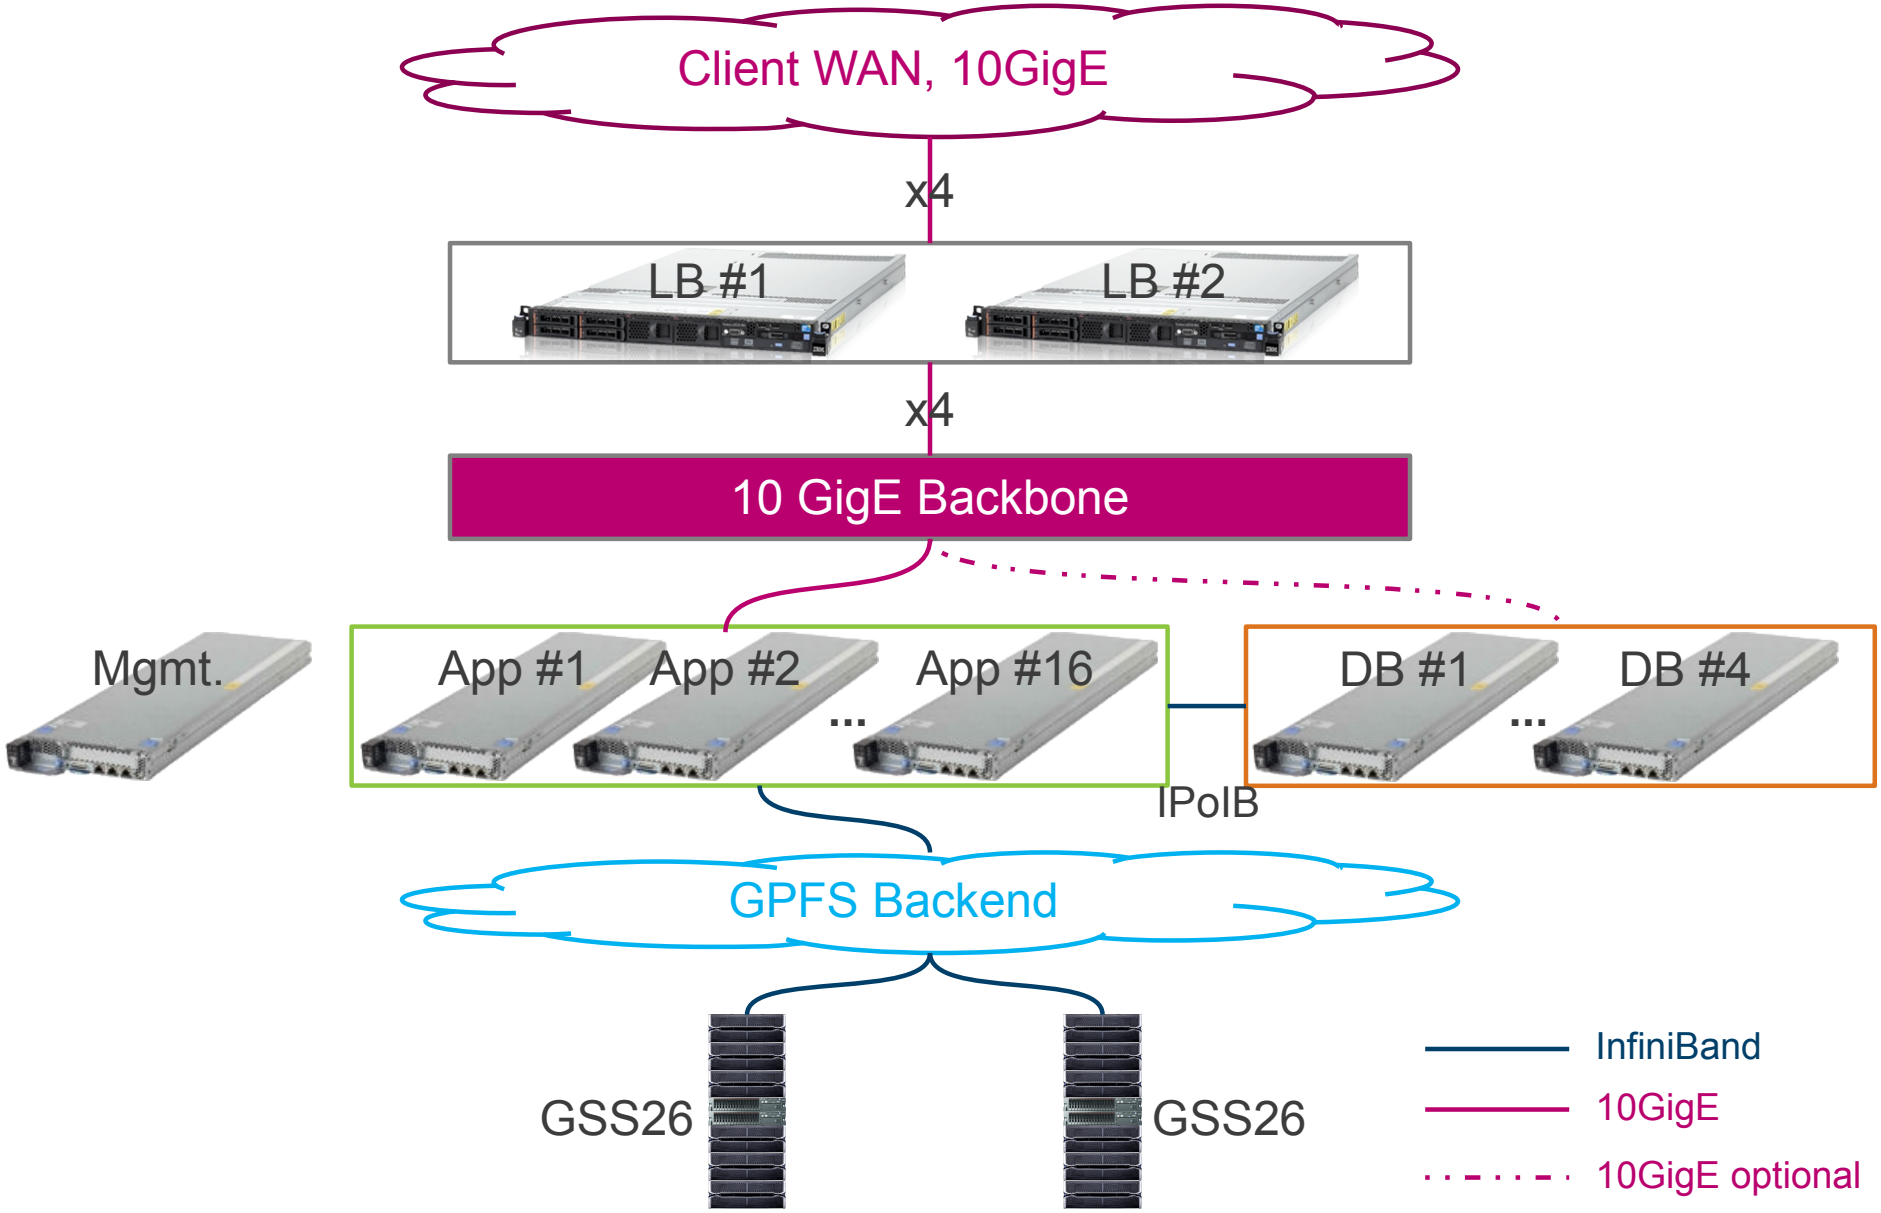

Software

- RHEL for GSS
- CentOS preferred, but RHEL acceptable
  - RHEL academia licenses
  - RHEL 7 outlook?
- LVS as load balancer Software (http://www.linuxvirtualserver.org/)
- XCAT (http://sourceforge.net/p/xcat/wiki/Main_Page/) as systems management
- ownCloud Enterprise Edition
- Nagios/Icinga for monitoring
- LDAP authentication

Hardware

- NeXtScale n1200 Chassis
  - 6 x 900W PSU
- Mellanox SX6036 InfiniBand Switch
- IBM G8264 10GigE Switch
- LG-E ES-4052G Ethernet / Management Switch
- GSS26
  - 4 TB disks
  - InfiniBand backend
  - GPFS Data over RDMA

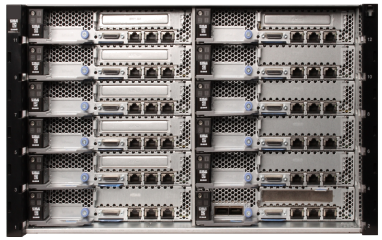

Network Concepts – Remarks

- App/DB nodes connected to 10GigE backbone
  - Optional, otherwise connect to client WAN directly
  - Optional 10GigE for DB node (due to use of IPoIB)
- App nodes connected to IB to
  - access GSS/GPFS
  - access DB nodes via IPoIB
    - DB nodes connected to IB by using IPoIB
- GPFS data over InfiniBand (RDMA)

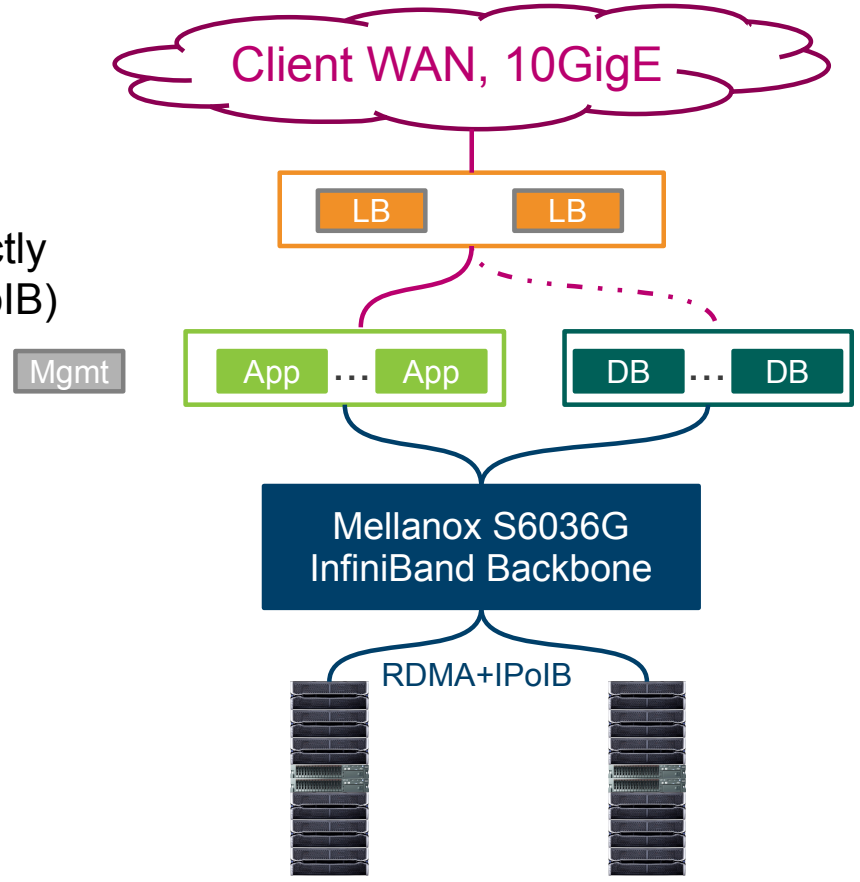

Summary and Outlook

- IBM Elastic Storage delivers perfect storage basis for ownCloud based sync & share cloud
- PoC delivered basic insights on sizings for huge installations and was very successful
- Great partnership with ownCloud → now official IBM Partner (see last slide)
- Business Partner Solution: pro-com Academic Private Cloud

- Next Steps:
  - Win more clients together
  - Deeper technical cooperation currently in discussion
  - Leverage more ownCloud Enterprise + GPFS capabilities (ILM) – see next slide
  - IBM Power-based ownCloud Infrastructures (LB/App/DB and Elastic Storage Server)
  - DB2 database backend in ownCloud
- Acknowledgements:
  - Frank Karlitschek (CTO, Founder) ownCloud
  - Oliver Kill and the pro-com team

Code-named Elastic Storage – unique solution feature richness

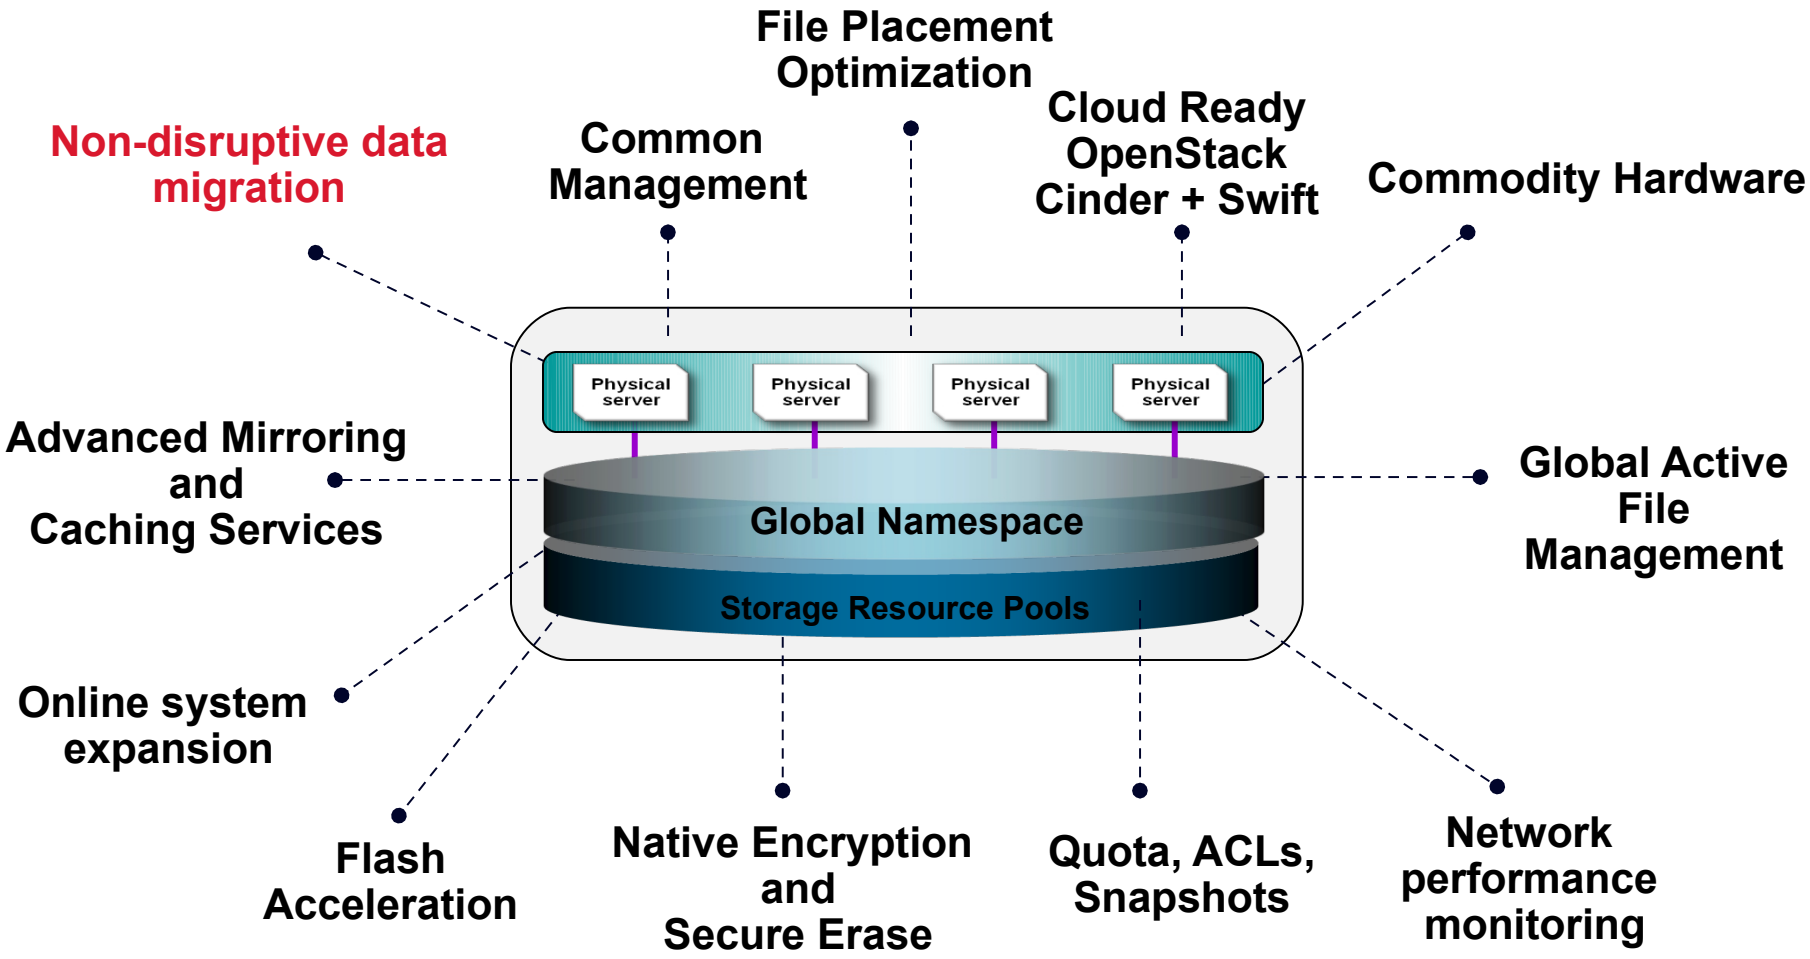

Fostering the fruitful partnership


IBM Systems > System Storage > Solutions >

IBM solutions from independent software vendors, partners and solution providers

Resource library | Quotes

"As organizations increasingly build on-premises, private clouds, IBM's enhancements to their Elastic Storage offerings provide an ideal data infrastructure for ownCloud's enterprise file sync and share (EFSS). A deployment of ownCloud Enterprise Edition requires an underlying storage platform such as IBM's Elastic Storage that is highly extensible and creates a reliable service for high availability. Elastic Storage removes data-related bottlenecks by providing parallel access to data, eliminating single filer choke points or hot spots and also simplifies data management at scale. We are excited to partner with IBM as they continue to leverage innovative technologies and software defined storage capabilities to enable faster information delivery, increased efficiency and scalable performance for cloud environments." -- Markus Rex, CEO and Founder

ownCloud is IBM Storage Preferred Partner

IBM Storage Preferred Partner Website

ownCloud is an IBM New Routes Partner

Agreement in place for NA, work in progress for EMEA & R-o-W

New Routes: "…is designed to help IBM and select IBM Business Partners capitalize on emerging growth opportunities, win new clients and increase market share as part of a joint marketing effort…."

IBM New Routes website

ownCloud and IBM to co-present at various events

Panel discussion at Fast Data Forum in Boston

ISV/MSP Mashup @ Enterprise2014

STU Dublin 2014

Edge 2015

more to come…

Trademarks and disclaimers

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: AIX®, alphaWorks®, BladeCenter®, Cognos®, Cool Blue™, DB2®, developerWorks®, Diligent®, DS4000®, DS6000™, DS8000®, Easy Tier™, Enterprise Storage Server®, ESCON®, eXtended I/O™, FICON®, FlashCopy®, GDPS®, Geographically Dispersed Parallel Sysplex™, GPFS™, HACMP™, HyperSwap®, IBM®, IBM TotalStorage®, IMS™, Lotus®, MVS™, Notes®, Parallel Sysplex®, POWER®, POWER7™, PowerHA™, ProtecTIER®, Rational®, Redbooks®, RMF™, Storwize®, System i™, System p™, System x™, System z™, System Storage™, System Storage DS™, Tivoli®, Tivoli Storage Manager Fastback™, TotalStorage®, WebSphere®, XIV®.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to: IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

The following are trademarks or registered trademarks of other companies:

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

ownCloud and the ownCloud Logo is a registered trademark of ownCloud, Inc. in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Snapshot, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.UNIX is a registered trademark of The Open Group in the United States and other countries or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

SPC Benchmark 1, SPC-1, SPC-1 IOPS, SPC-1 LRT, SPC Benchmark 1C, SPC-1C, SPC Benchmark 1C/Energy, SPC-1C/E, SPC Benchmark 2, SPC-2, SPC Benchmark 2C, SPC-2C, SPC Benchmark 3BR, and SPC-3BR are trademarks of the Storage Performance Council.

Other company, product, or service names may be trademarks or service marks of others.

Trademarks and disclaimers (cont.)

NOTES:

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area. Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography. Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

Information is provided “AS IS” without warranty of any kind. The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services.