

Open Enterprise FileStore



Stonehenge: "virtual distributed NAS"



Stonehenge Marketing
stonehenge@de.ibm.com

IBM TotalStorage™ Connected. Protected. Complete.

© Copyright IBM Corporation 2003. All rights reserved.

Stonehenge at a Glance

- Enterprise NAS Gateway for NT, Unix, Linux, AFS
- Virtualized Storage
- OpenSource- & Unix-based
- One Administrator per Geography
- One System per Company

Stonehenge at a Glance

- 100.000+ clients (users)
- 1 Petabyte+ hosted data
- Maximum Redundancy
- Maximum Security

- Web-/Serviceprovider friendly
- DCE-DFS replacement

Why Stonehenge was developed

- Needed **24h×7** storage environment for disk plant
- Needed to include numerous unrelated divisions
- Needed to attach **foreign** production **sites**
- Needed easy **scalability** due to unknown growth
- Needed online **migration** from running environments

- Didn't want to purchase more MS **licences**
- Didn't want to service non-consolidated hardware

What Mainz implemented

- 3000 Users file storage, Multi-Terabyte
- Distributed System, accessible from Hungary & California
- Mature core components: Linux, AIX, OpenAFS
- Consolidated ~100 NT servers into 2 racks (+ SAN disks)

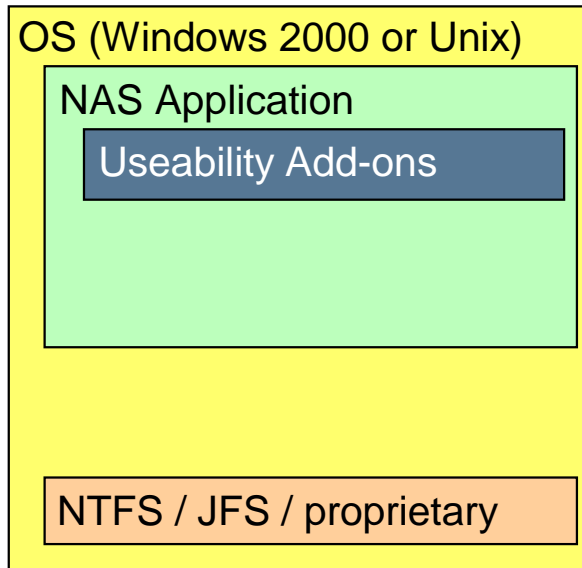
- Savings in 2002:



- Licences 300000 \$
- Administration 3-4 Headcounts
- No Downtime ??? \$

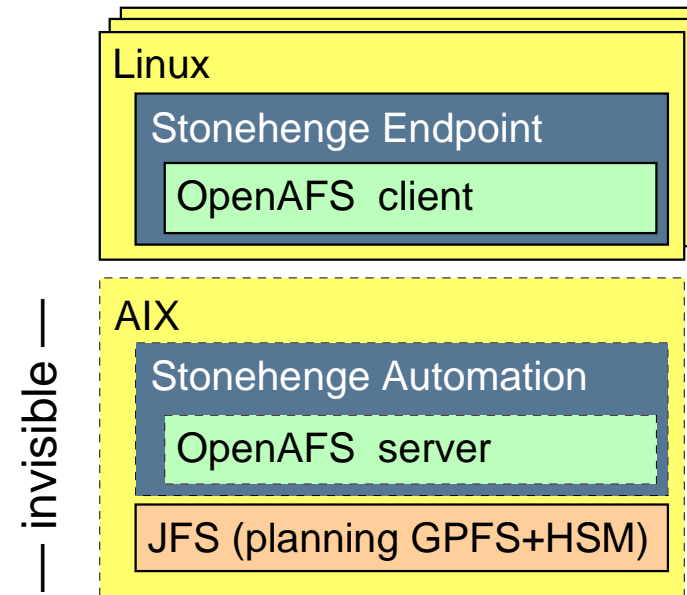
Comparing with IBM NAS Gateway

- NAS Gateway



SAN Storage


- Stonehenge



SAN or 7133-SSA

Server Virtualization included


- Virtual file servers across the Linux farm
 - set up in minutes
 - auto-restarted in case of crash
 - auto-rebuilt in case of failing hardware



- Benefits:
 - Maximum redundancy, "always on"
 - Easy scaling
 - No Maintenance windows needed
 - User/Client believes seeing a standard Windows or Unix fileserver

Subsystem Virtualization included

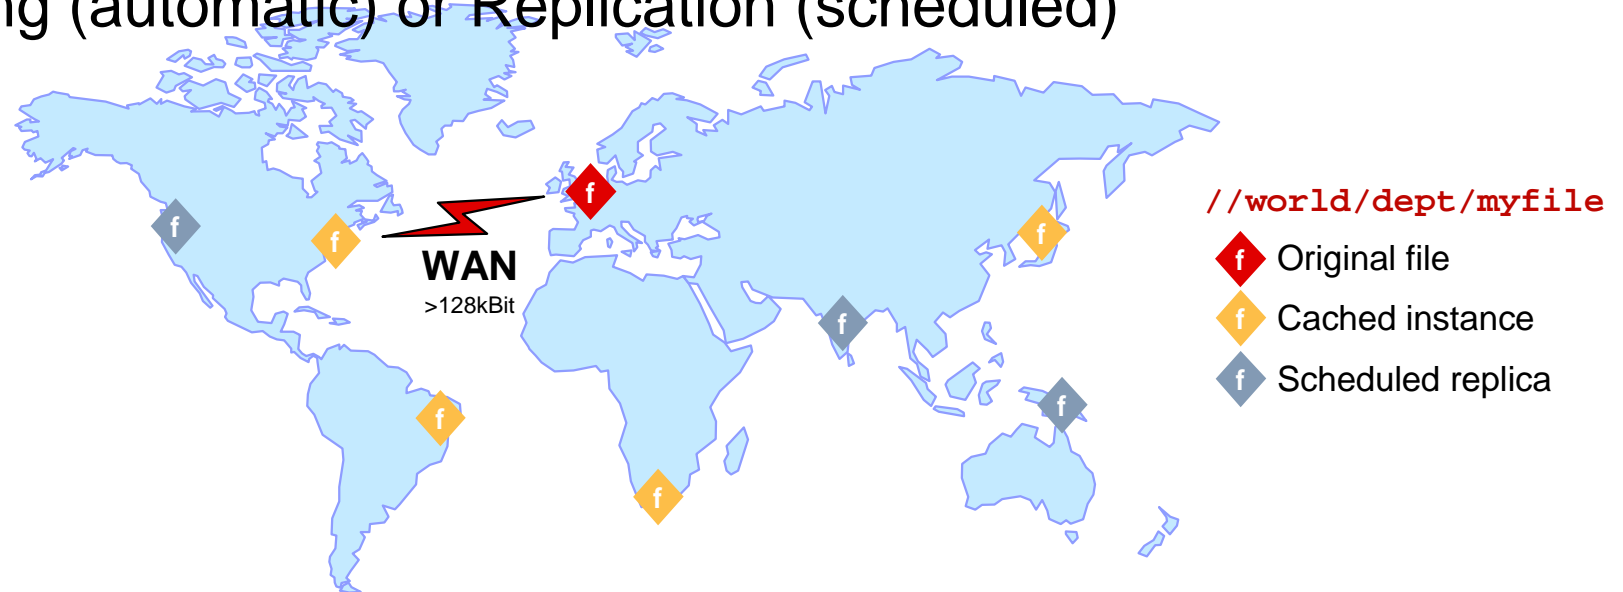
- Storage is transparent
 - Migrate storage-to-storage without service interruption
 - Define storage space across several subsystems



- Benefits:
 - No maintenance window for storage addition / migration / servicing
 - Optimum usage of latest storage hardware
 - Long-distance (1000km+) async. replication included

One System per Company

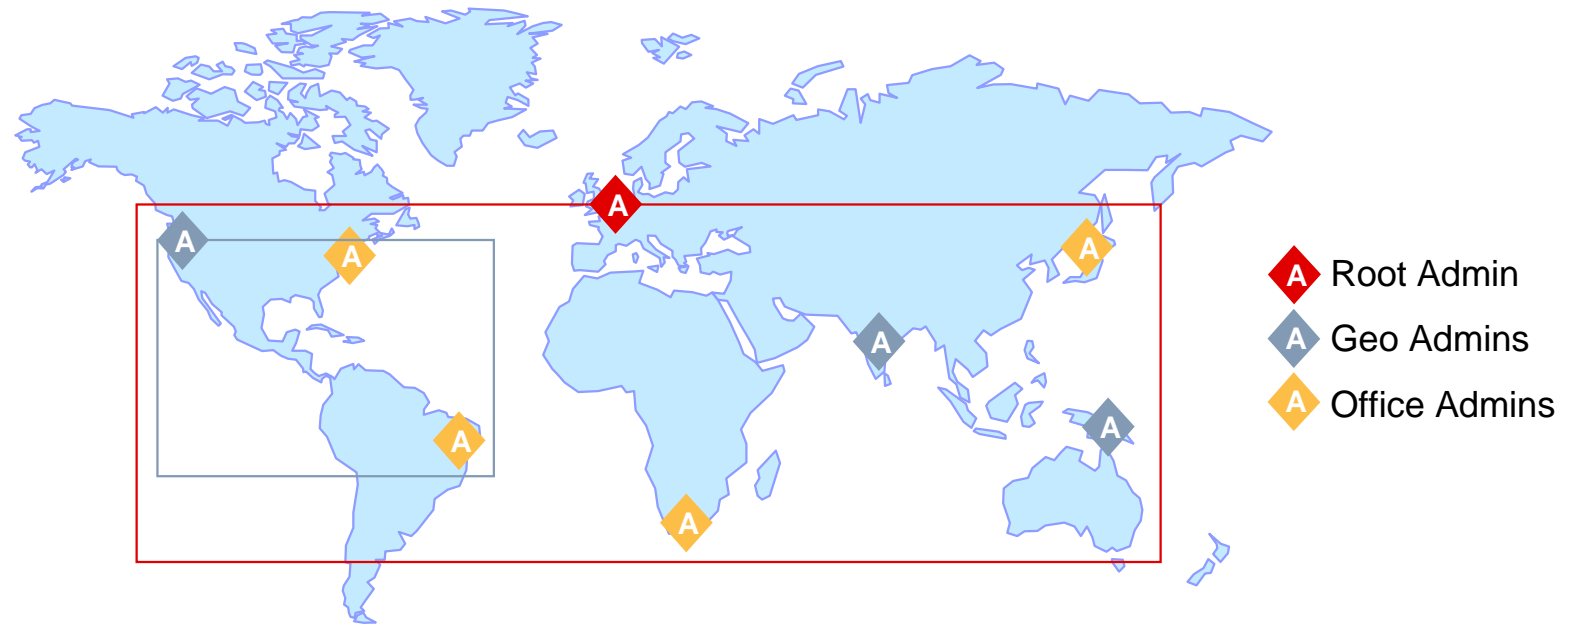
- Scale one system accross continents
- Access anything from anywhere
- Caching (automatic) or Replication (scheduled)



- **One file tree**, always serving nearest instance of a file

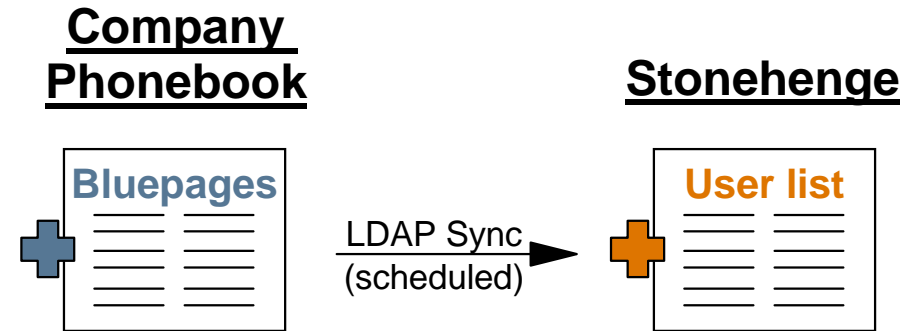
One Administrator per Geography

- Administration is hierarchical
- Root Admin has Company-wide view




Intelligent Automation


- IBM Employee newly hired:



- IBM Employee re-assigned:



- IBM Employee retires:



- **Minimum Admin involvement !**

Intelligent Automation

- Automated Backup/Restore including all access rights*
- Migrate data between storage subsystems while online
- Quota etc.

| AFS Volumes | | | | | | |
|------------------|-----------|-------------|-----------|-----------|----------|--------|
| Name | Partition | Server | Type | Used | Quota | Used % |
| labor | /vicepc | mfgmzafsc01 | 536905544 | 10 KB | 1,02 MB | 1 % |
| install_w2k_s... | /vicepc | mfgmzafsc01 | 536870966 | 265,91 MB | 409,6 MB | 64 % |
| sas_mzhdd_q... | /vicepc | mfgmzafsc01 | 536871541 | 14,99 GB | 51,2 GB | 29 % |
| dept_2128 | /vicepc | mfgmzafsc01 | 536871190 | 85 KB | 5,12 GB | 0 % |
| dept_4328 | /vicepc | mfgmzafsc01 | 536871196 | 35,33 GB | 71,68 GB | 49 % |
| dept_4603 | /vicepc | mfgmzafsc01 | 536871202 | 6 KB | 5,12 GB | 0 % |
| dept_4648 | /vicepc | mfgmzafsc01 | 536871205 | 6 KB | 5,12 GB | 0 % |
| cvs | /vicepc | mfgmzafsc01 | 536871154 | 475,37 MB | 1,02 GB | 46 % |
| dept_4655 | /vicepc | mfgmzafsc01 | 536871211 | 2 KB | 5,12 GB | 0 % |
| rsca_bdw | /vicepc | mfgmzafsc01 | 536871654 | 968,55 MB | 2,05 GB | 47 % |
| install | /vicepc | mfgmzafsc01 | 536870936 | 16 KB | 5,0 MB | 0 % |
| rsca_develop | /vicepc | mfgmzafsc01 | 536871657 | 6,83 GB | 10,24 GB | 66 % |
| dept_4119 | /vicepc | mfgmzafsc01 | 536871193 | 787,16 MB | 5,12 GB | 15 % |
| dept_4652 | /vicepc | mfgmzafsc01 | 536871208 | 525,51 MB | 5,12 GB | 10 % |
| install_xp_prof | /vicepc | mfgmzafsc01 | 536870975 | 475,65 MB | 1,02 GB | 46 % |
| install_xp | /vicepc | mfgmzafsc01 | 536870972 | 9 KB | 5,0 MB | 0 % |
| install_w2k_a... | /vicepc | mfgmzafsc01 | 536870963 | 266,95 MB | 614,4 MB | 43 % |
| install_apps | /vicepc | mfgmzafsc01 | 536870939 | 6 KB | 5,0 MB | 0 % |
| install_apps_... | /vicepc | mfgmzafsc01 | 536870945 | 3 KB | 5,0 MB | 0 % |
| install_logs | /vicepc | mfgmzafsc01 | 536870951 | 10 KB | 102,4 MB | 0 % |

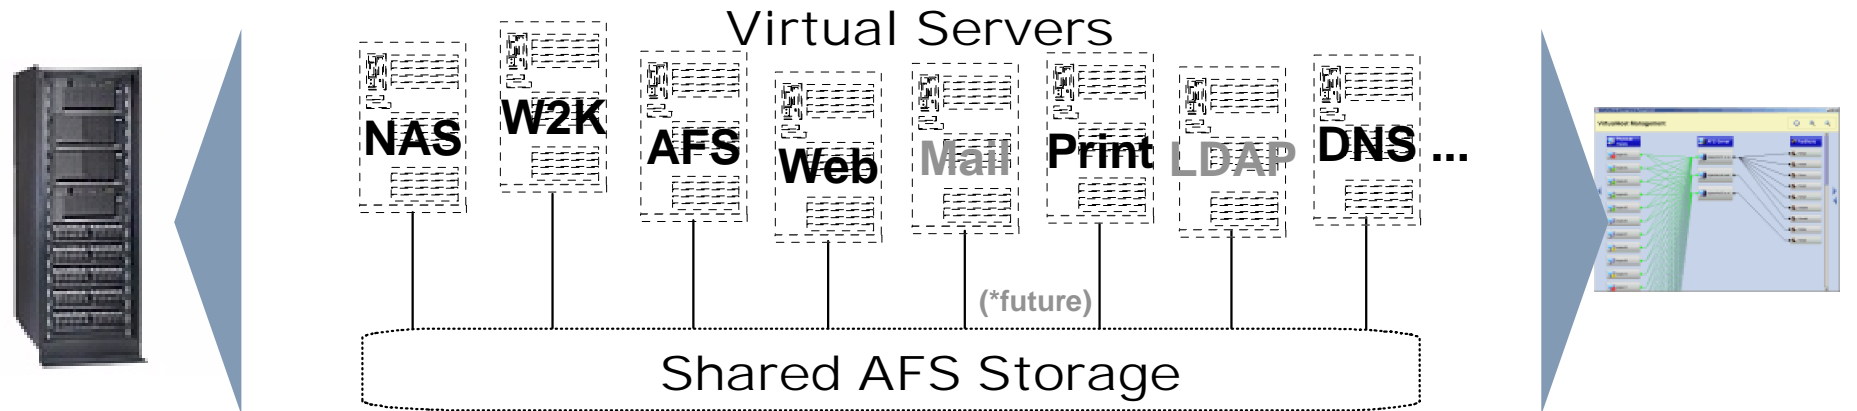


- Set Quota
- Release
- Salvage
- Backup
- Move
- Mount
- Rename
- Delete
- Lock
- Unlock

*prerequisite: TSM


Options/Plugins

- Separate web access rights (Filesystem = Intranet)
- Synchronization with existing LDAP service e.g. IBM Bluepages
- Integrated Management of Active Directory®
- Integrated Management of MS dfs
- Integrated Tivoli Storage Manager automation
- User-accessible Snapshot of yesterday's files
- Online-Virusscan during file transfer
- Infrastructure (DHCP, DDNS,...) serviced by Stonehenge

IT Infrastructure Transformation




1 - Enterprise Solution ~3000 users




- ✓ Scalable Fileservice
- ✓ Transparent failover
- ✓ Windows/Unix/AFS/ Web clients
- ✓ Single management interface (Java)

2 - Split Sites Solution



- ✓ One Namespace
- ✓ With/without (*) Disaster Recovery
- ✓ With/without async. replication
- ✓ Site distance >>100km

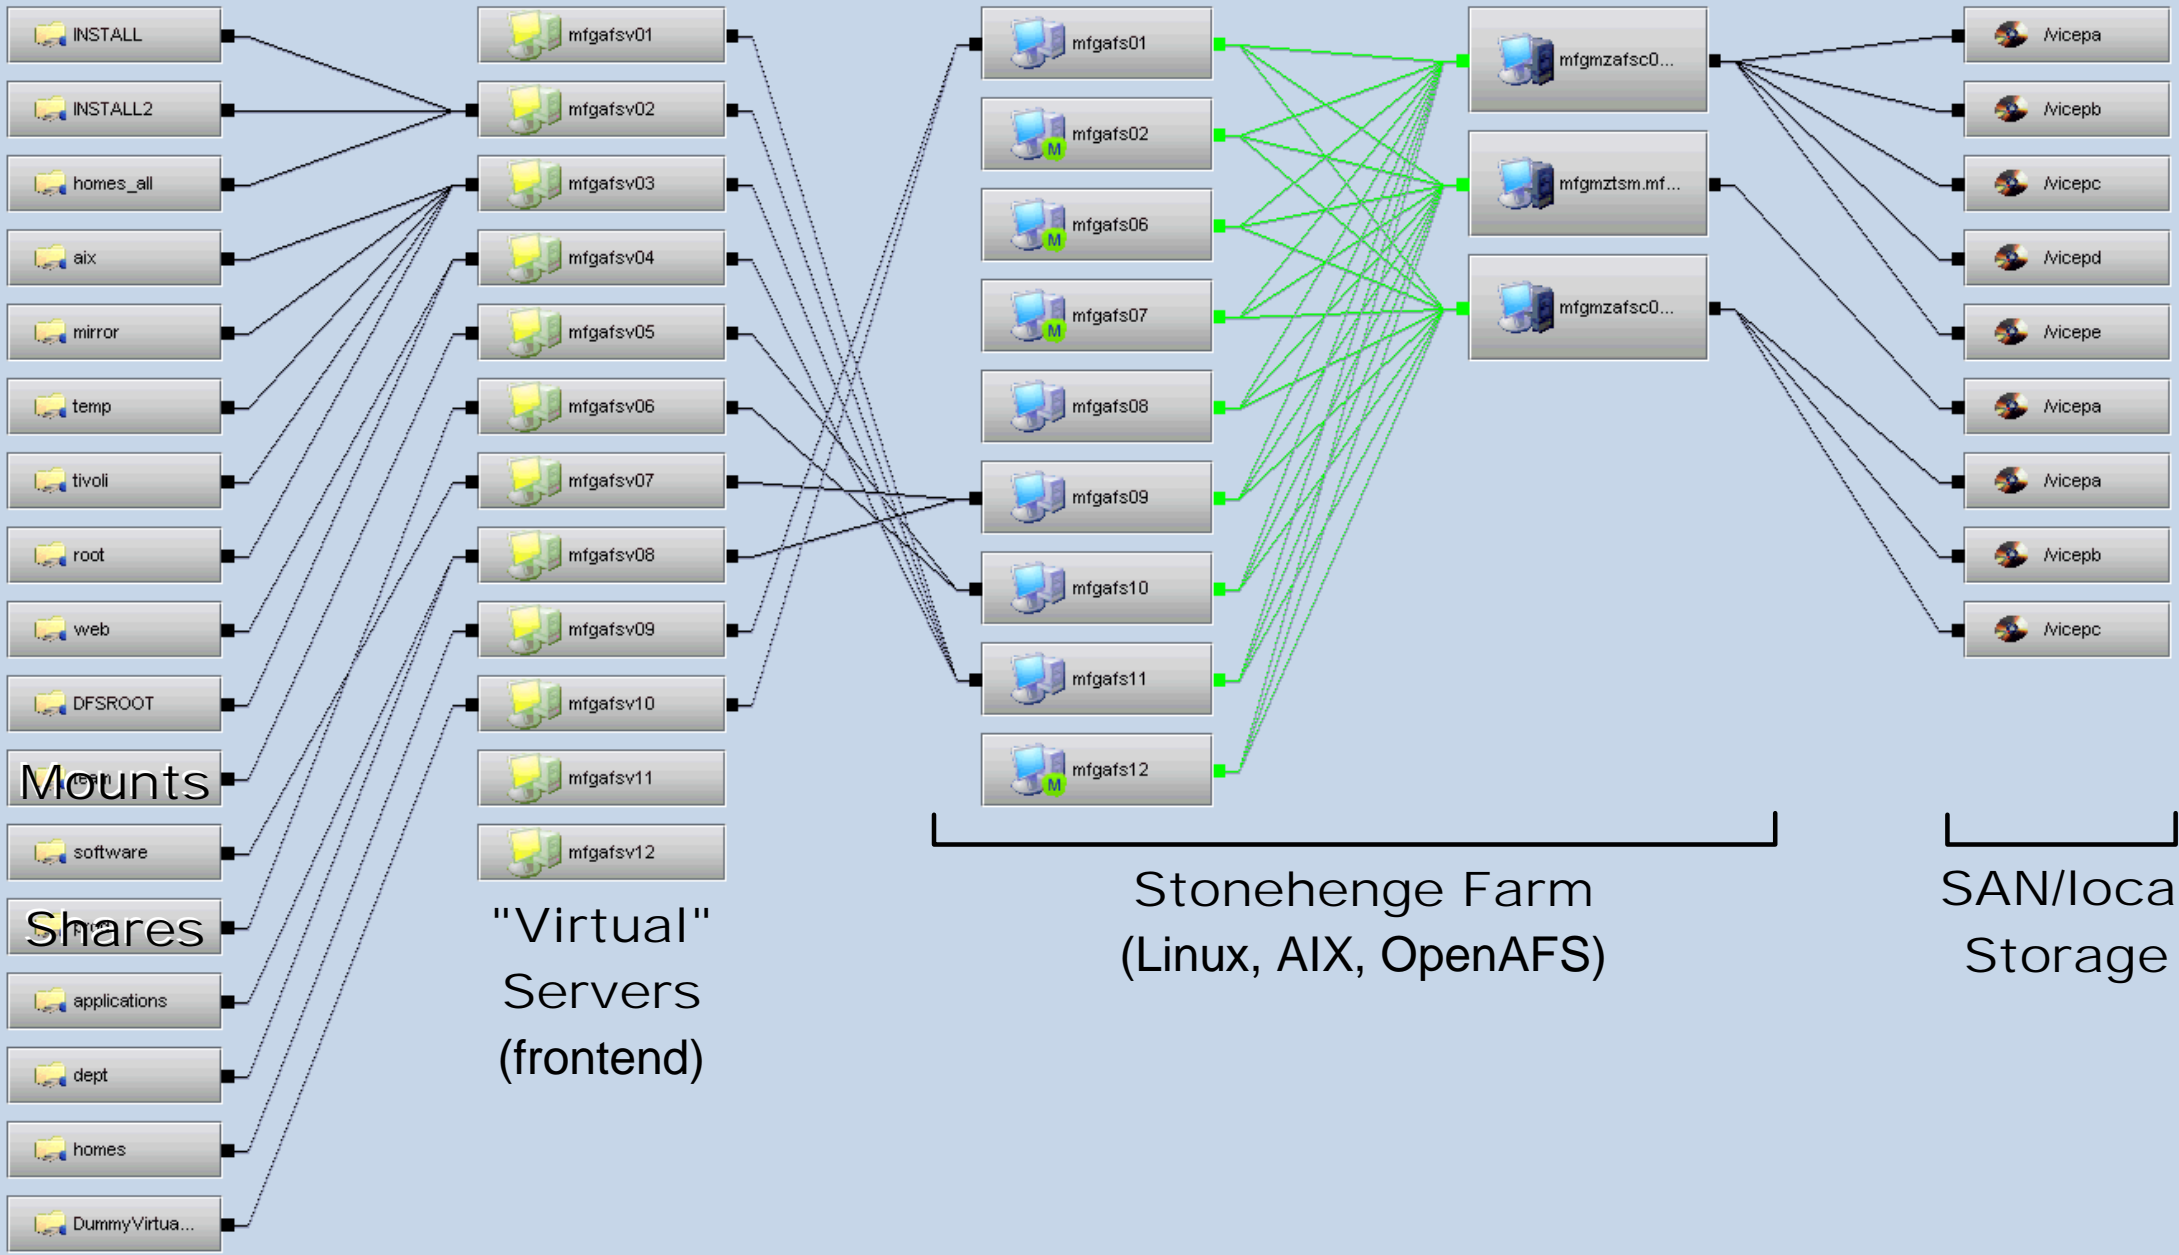

3 - Head & Branch Offices Solution



- ✓ True Filesharing
- ✓ Intelligent Caching
- ✓ Branch offices >>100km away

Administrator Screenshot: Virtual Server distribution

VirtualHost Management



Administrator Screenshot: Load Balancing

Monitoring & Statistics

Menu

Network monitoring ▾

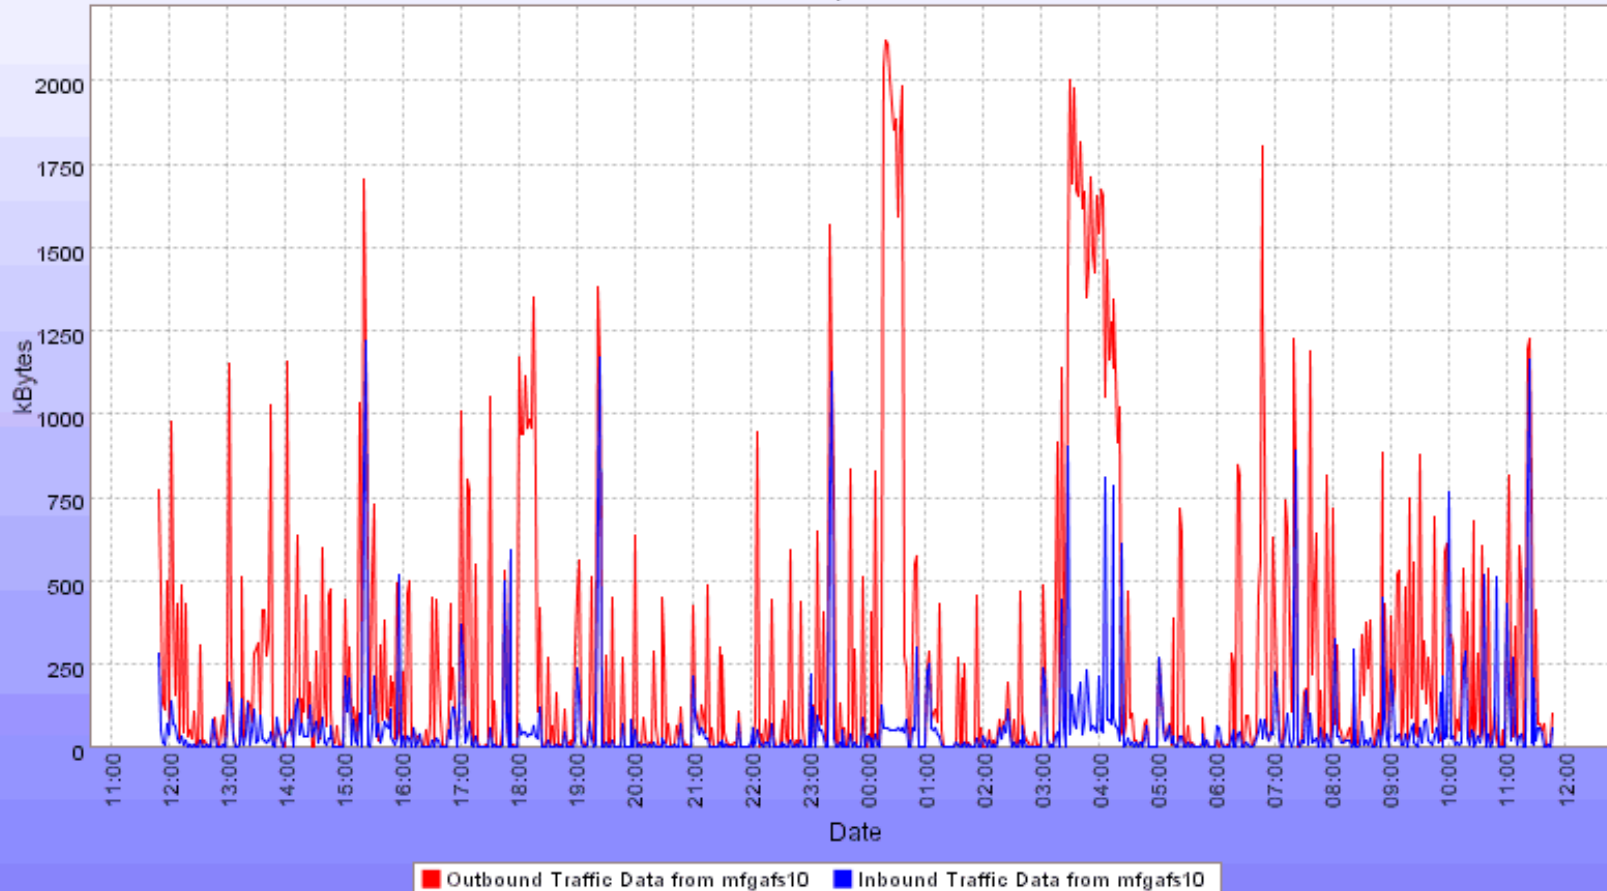
Daily Weekly
 Monthly Yearly

Traffic inbound
 Traffic outbound



Average per minute in kBytes from mfgafs10

Daily



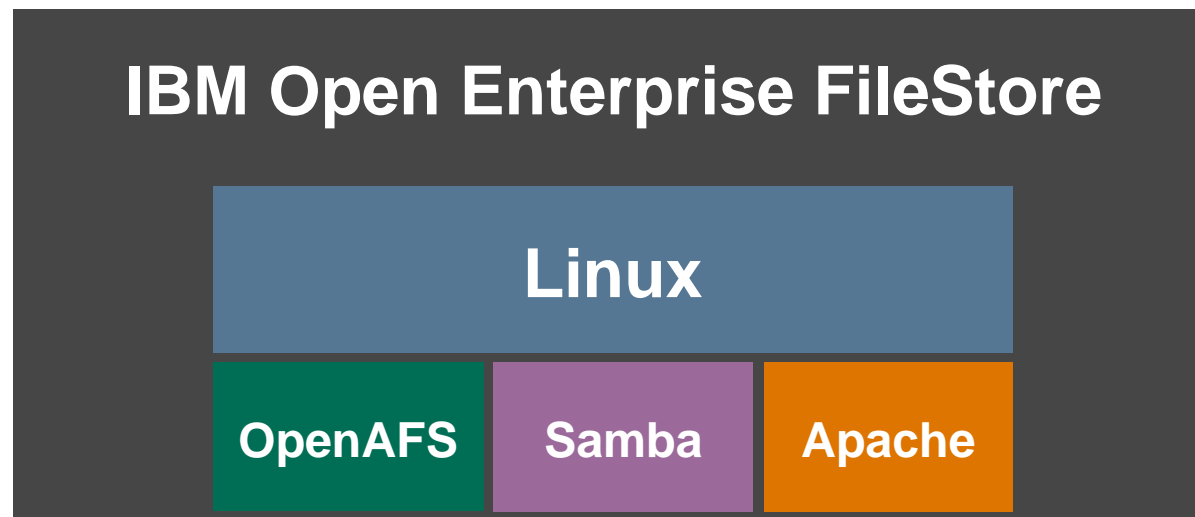
Property of IBM Corporation, in the United States, other countries, or both.

t:mfgafs10

User: Mfg-Mainz\de118196

Open Standards

- No network licences for attached Windows environment
- **Linux** Distributions used Redhat & Suse
- Samba 2.2.8a, Apache 2.x/1.3.x, OpenAFS 1.2.8
- Kerberos 5 security & encryption



Disclaimers

The sizing data contained herein was obtained in a controlled environment based on the use of specific data. Actual results that may be obtained in other operating environments may vary significantly. These values do not constitute a guarantee of performance.

Product data is accurate as of initial publication and is subject to change without notice.

No part of this presentation may be reproduced or transmitted in any form without written permission from IBM Corporation.

References in this document to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

The information provided in this document has not been submitted to any formal IBM test and is distributed "As Is" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into their operating environment.

While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.