



# *Samba in Enterprise Environments*

e**XP**eriences and problems encountered

Joachim Jordan  
Michael Weisbach  
IBM Linux Competence Center Boeblingen

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*	ESCON*	Netfinity	System/390*
CICS*	FICON	OS/390*	VM/ESA*
DB2*	IBM*	PR/SM	VSE/ESA
DB2Connect	IBM logo*	RS/6000*	Virtual Image Facility
DB2 Universal Database	IMS/ESA	S/390*	VisualAge*
e-business logo	MQSeries*	S/390 Parallel Enterprise Server	WebSphere
	Multiprise*	SecureWay	zSeries

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation

Tivoli is a trademark of Tivoli Systems Inc.

LINUX is a registered trademark of Linus Torvalds.

Penguin (Tux) compliments of Larry Ewing.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

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.

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.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

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.

## ***Agenda***

- **Overview over some customers**
- **A customer in detail**
- **Features requested**
- **Discussion: Samba in large enterprise environments**

## ***Some customers***

- **German power supplier:**
  - ▶ 110 servers, ~40,000 user IDs via Winbind, 3.5 TB data
- **German parliament:**
  - ▶ ~1.500 Clients, about 100 decentralized server
- **German insurance company:**
  - ▶ ~1.000 servers
- **Danish bank:**
  - ▶ ~600 servers
- **US insurance company:**
  - ▶ 12 servers, 3.6 TB data

## ***A customer in detail - ongoing Proof of concept***

### **■ German power supplier:**

- ▶ currently using Windows NT servers and clients
- ▶ 110 file servers in main location
- ▶ in summary ~3.5 TB data on these servers; SCSI disks
- ▶ ~30.000 active clients

### **■ now evaluating**

- ▶ Fileserving with Samba on Linux for S/390 vs. W2k/WinXP

### **■ future**

- ▶ Active Directory infrastructure
- ▶ Samba/390 *and/or* W2k/Windows XP based file servers
- ▶ Windows XP clients

## ***Summary experiences***

- **easy administration of repetitive tasks**
- **large user databases drive winbind requirements**
- **large storage requirements and clustering issues**
- **high failure costs drive need for redundancy and replication in all major components**
- **large enterprises need to lower admin costs, especially in heterogenous environment - (Active) Directory coexistence/integration**
- **Enterprise wide security via addition of Kerberos**
- **large enterprises are multi-national - internationalization requirement**

## ***Features needed - Access Control***

- **migrations from non-UNIX to UNIX**
  - ▶ from OS/2, Novell Netware and Windows based servers
  - ▶ to Samba on Linux (UNIX)
  
- **ACL support to match "old" access controls**
  - ▶ journaling filesystem required
    - ◆ possible with XFS or ext3 with ACL patches
    - ◆ JFS ACL support planned
    - ◆ not official supported yet (kernel 2.4.x nor any current distributions)
  - ▶ POSIX ACLs are slightly different to Windows implementation
    - ◆ missing "append only" and "write/edit but not delete"
    - ◆ Solution: EA based access control vs. Samba-internal ACL database?
  - ▶ Problem
    - ◆ different access paths: Samba or via UNIX shell
    - ◆ ACL performance for backup/recovery?

## ***Features needed***

- **Winbind mapping is not unique**
  - ▶ today dynamic allocated
  - ▶ Backup is done for Unix ID, how to map correctly to NT RID in case of restore?
  - ▶ missing import/export; ASCII, xml?
  - ▶ unique mapping also needed for
    - ◆ replication/mirroring for high availability
    - ◆ scalability via clustering and loadbalancing
  
- **Backup**
  - ▶ backup tools need to backup all necessary data
    - ◆ ACL information
    - ◆ winbind mappings
  
- **Virus scanner integration**



## ***Features needed - Auditing and Control***

### ■ **expectations**

- ▶ customers w/ WindowsNT servers using (and expecting) auditing features
- ▶ n:1 migration (integration) of file servers requires detailed accounting

### ■ **Logging, Auditing and Accounting**

- ▶ functions to support logging and auditing of user activity
- ▶ accounting
  - ◆ Samba process runs under security of user, therefore process accounting may be used
- ▶ no text logfiles please :) - API layer preferred

## ***Features needed - UserAdmin for LDAP***

- **challenge: German Parliament decided use OpenSource on all servers; not to use Active Directory but instead native LDAP!**
  - ▶ large, decentralized Samba environment on Intel based servers
  - ▶ native LDAP directory
  - ▶ Windows XP clients
  
- **problem: How to administer, create and delete users?**
  - ▶ use of NT tool "Domain Administration for Domains"
    - ◆ how to assign UNIX UID to new LDAP (sambaAccount) user?
  - ▶ have a samba account admin frontend w/ LDAP support
  - ▶ ...

## ***Features needed - High Availability***

- **High Availability and Scaleability needs**
  - ▶ support for shared (or replicated) file systems
    - ◆ not only VFS support
    - ◆ maybe cluster filesystem (GFS, GPFS, ...)
    - ◆ use of OpenAFS filesystem across many file servers (i.e. for load balancing)
  
- **unique Winbindd - NT username to UNIX uid mapping required**
  - ▶ two or more nodes of failover cluster / high availability
  - ▶ two or more nodes of load balancing cluster / more than one path to file

## ***(Known) Limitations***

- **User ID's**
  - ▶ currently only ~65000 userids /groupids possible
  
- **Sockets: TCP/IP stack limitations per Linux instance**
  - ▶ solution: "Samba farm" w/ shared filesystem?!
  - ▶ inside Samba server box, additional sockets needed too (LDAP connects)
  
- **Processes / Linux box limited**
  - ▶ Linux process scheduler scalability
  - ▶ multiprocess / multithreaded approach like Apache 2.x?
  
- **Memory limitations**
  - ▶ solution: 64bit support

## **Contact**

- **Michael Weisbach, Linux IT architect,  
IBM Germany Global Services**
  - ▶ Linux IBM IGS Germany: [crlinux@ch.ibm.com](mailto:crlinux@ch.ibm.com)
  
- **Joachim Jordan, Linux IT architect,  
IBM Europe technical sales support**
  - ▶ Europe Linux ATS: [tmcc@de.ibm.com](mailto:tmcc@de.ibm.com)
  
- **<http://www.ibm.com/linux>**

## *Discussion*

- **Feel free to comment...**