

What's new on SMB Traffic Analyzer

A look back to 2009, where we are currently, and a short glance into the future.



Holger Hetterich
Level 3 technical support engineer
SUSE Linux Products GmbH

Benjamin Brunner
Trainee
SUSE Linux Products GmbH

Overview

- What is SMB Traffic Analyzer?
 - Short introduction and overview
- What happened after 2009 till today?
 - Taking a look at the past
- Current state
 - Introducing the client programs and the web interface
- Where do we go to, combined with Q&A

What is SMB Traffic Analyzer?



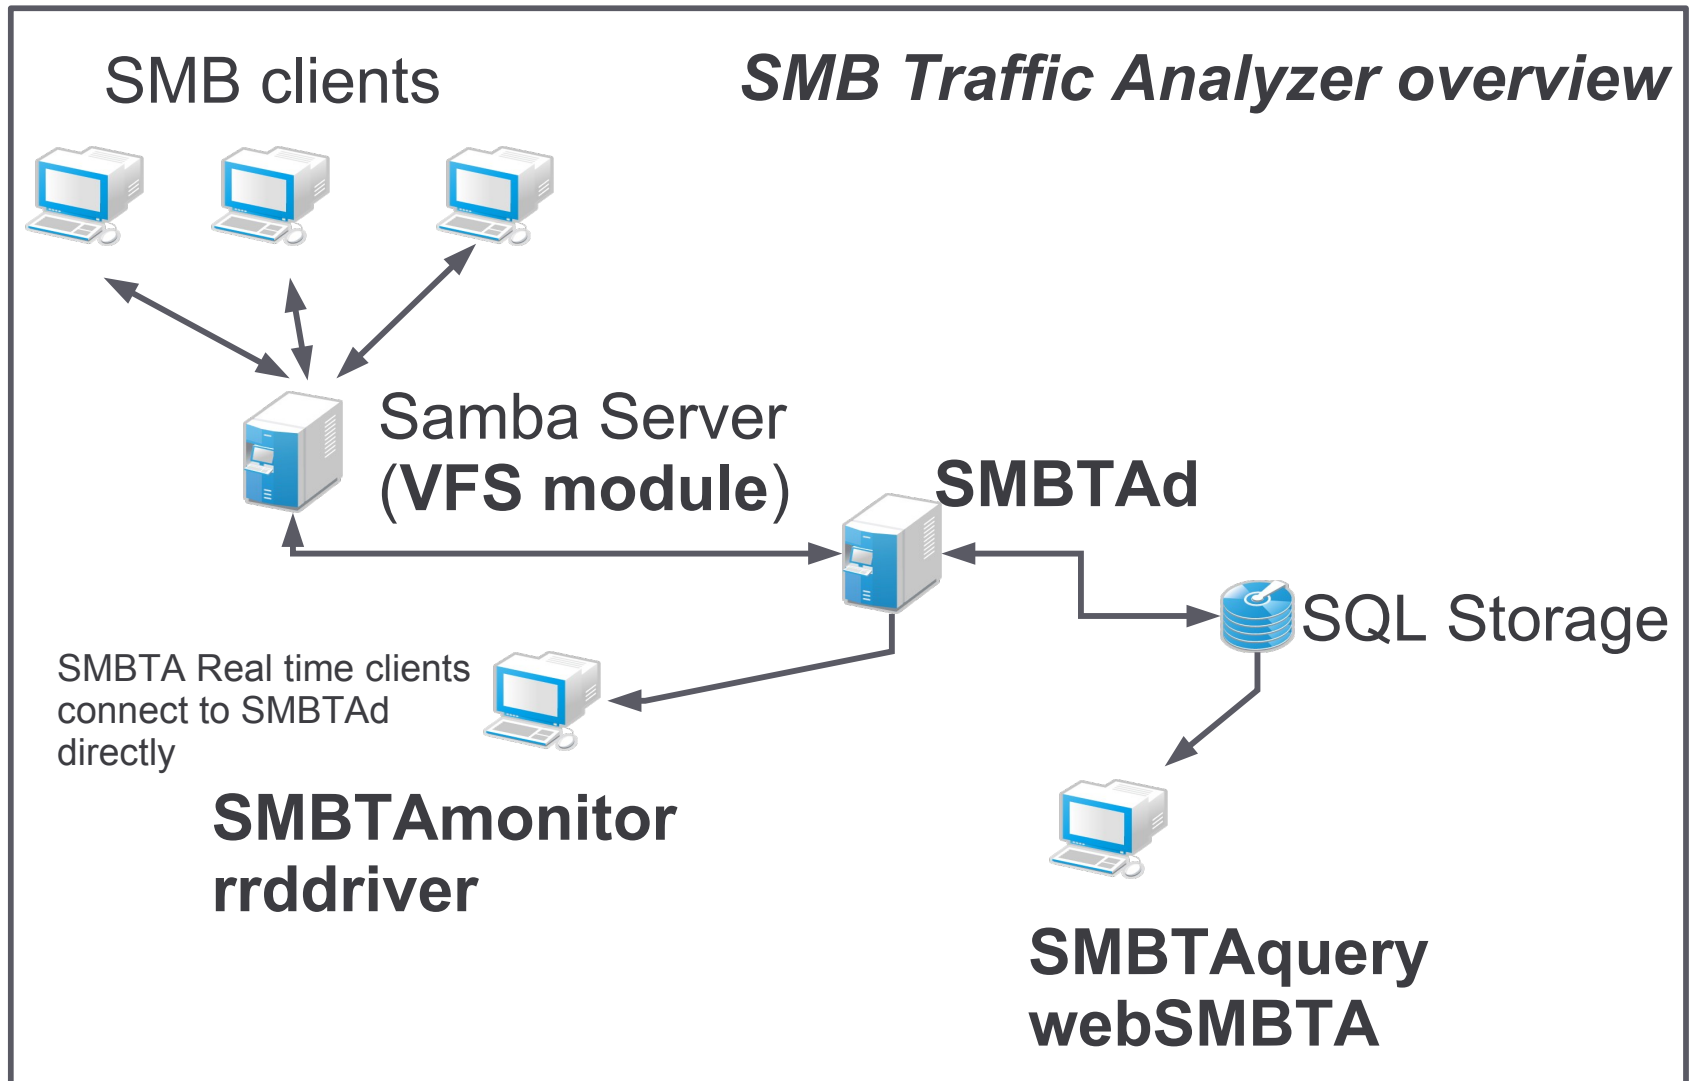
What is SMB Traffic Analyzer?

- SMB Traffic Analyzer (in short SMBTA) is a project aiming for answers to questions like the following:
 - What are the most used shares in my Samba network?
 - Which file is almost never used?
 - Which users are the most pressing in my Samba network?
 - At which time of the day, my Domain has the highest throughput?
 - Can I real-time monitor what happens on one of my shares?

What is SMB Traffic Analyzer?

- A Samba VFS (Virtual File System) module (`vfs_smb_traffic_analyzer`) to capture meta data of VFS operations.
 - For example: sizes of read and write operations
- Transfer this data to a receiver program (SMBTAd), which creates a SQL queryable database of it, or allow real time monitoring of the data flow
- End user tools to make it easy to query the database
 - SMBTAmoitor – real time monitor the data flow
 - SMBTAquery – produce statistics and charts in different output formats
 - WebSMBTA – intuitive web interface
 - Rrddriver – create rrdtool databases and graphs in real-time

SMB Traffic Analyzer concept overview



What happened since 2009?



What happened since 2009?

- SMBTA project milestones:
 - April: Tech-Demo in 2009 at SambaXP
 - > Was highly welcomed by users and developers but had problems
 - » No encryption of data
 - » Limited VFS function set that is parsed
 - » And more...
 - Summer: Thinking about concepts on how to go on and partying.

What happened since 2009?

- SMBTA project milestones in 2009:
 - 8th of October: *Madita was born*, a long pause in development
 - “If you think two kids are easy to handle, get a third !!”



What happened since 2009?

- Early 2010:
 - Base work for a completely rewritten VFS module was done and accepted by the Samba team:
 - > New, flexible data transfer protocol (v2)
 - > Supports encryption
 - > Ready for being extended (maybe support compression in future)
 - > Handle more VFS function types (such as rename, chdir etc)
 - > Store more basic data, like the SIDs of users or IP addresses of clients accessing the samba server

What happened since 2009?

- 2010 continued....
 - Middle of the year 2010:
 - > Base work for a toolset “smbttools” and the data receiver “SMBTAd” that supports the new VFS module
 - September 21: Presentation at SNIA SDC conference in Santa Clara, CA
 - October to November : SMB Traffic Analyzer 1.0 – 1.2.1
 - > Created SMBTAmonitor, rddriver
 - > Build on Solaris

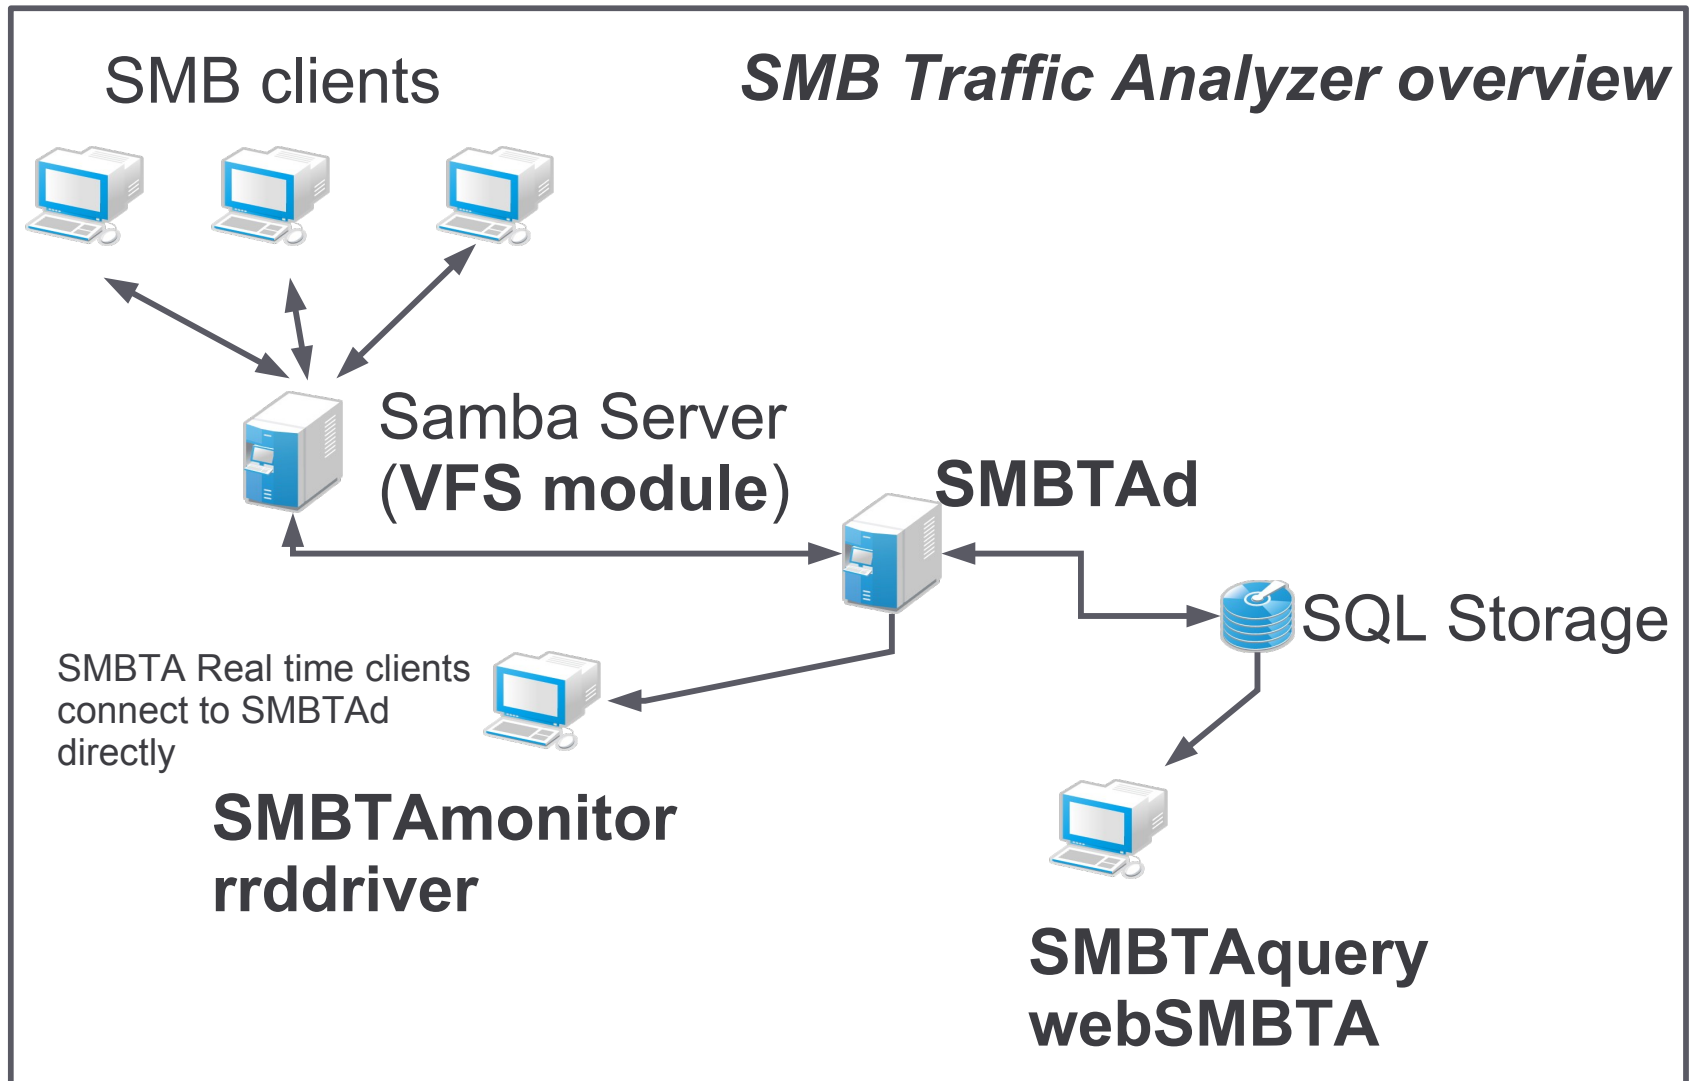
What happened since 2009?

- 2011:
 - January 20: released SMB Traffic Analyzer 1.2.2
 - > Major changes as we moved to XML for smbtaquery
 - » HTML and ascii output implemented
 - March 23, release SMB Traffic Analyzer 1.2.3
 - > Simplified and optimized smbtd to generate smaller databases and be much faster
 - May 08, SMB Traffic Analyzer 1.2.4
 - > Switch to libDBI, generic database interface

Current status



SMB Traffic Analyzer concept overview



Current Status – VFS module

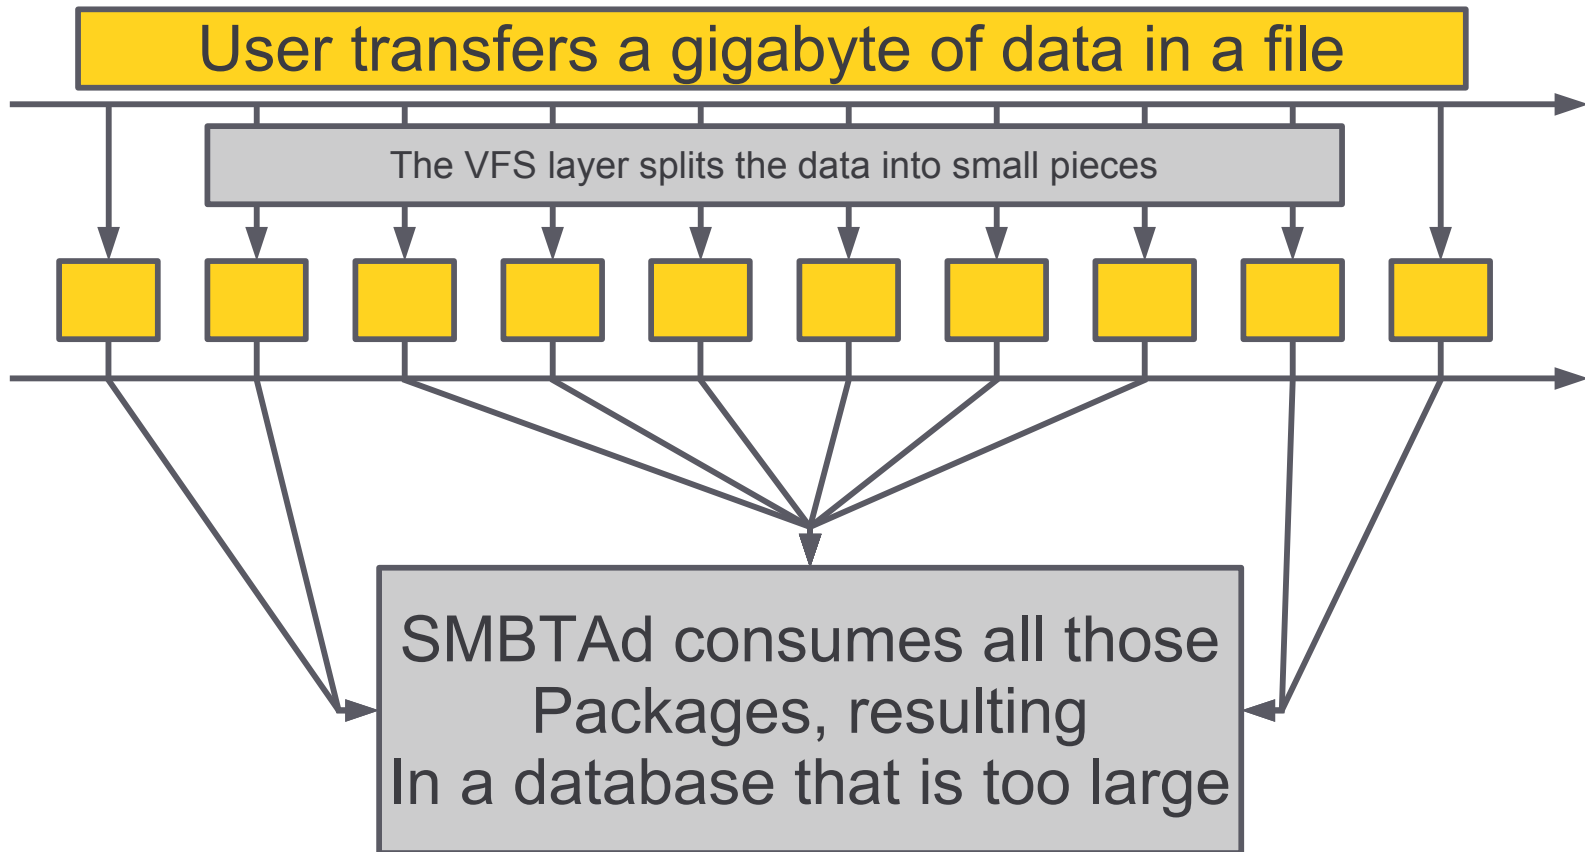
- Implements Protocol version 2
- Works on VFS functions: read, write, pread, pwrite, sendfile, chdir, rmdir, rename, close, open
- Is stackable, other VFS modules might be included while in operation
- IP-Address of the client machine and SID of the user are stored

Current Status -SMBTAd

- Bi-tree based, high performance cache algorithm, creating an interpolation of what happened through a user given timespan.
- Serving the database by libDBI, able to access and manage PostgreSQL, MySQL, Oracle, and sqlite


Current Status - SMBTAd

SMBTAd caching algorithm, initial problem




Current status - SMBTAd

SMBTAd caching algorithm, solution



Bi-Tree sort for the
Involved file...



... and only sum up
similar entries

Do this for 5 seconds, then flush the tree content into
the database.

Current Status - SMBTAquery

- Our main end user tool on the command line, a swiss army knife for SMBTA
 - Eases getting information about traffic flow by implementing a simple interpreted language:
 - > “global, total rw;”
 - Runs it's functions over Objects:
 - > Domains, Shares, Users, Files, Global
 - Creates XML and comes with stylesheets to automatically output to HTML or ascii

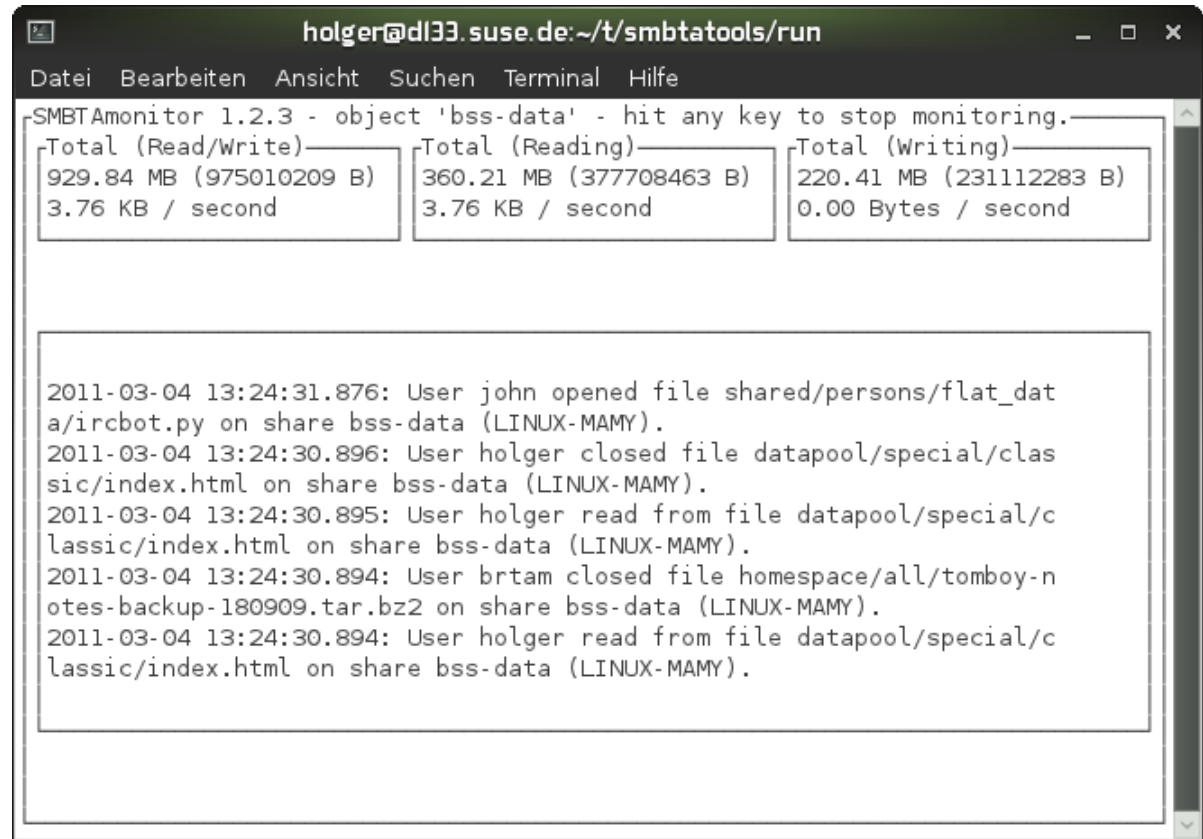
2:00 - 13:00	4.00 MB
3:00 - 14:00	19.37 MB
4:00 - 15:00	18.82 MB
5:00 - 16:00	60.54 MB
6:00 - 17:00	15.22 MB

Top 10 files globally by read-wri

Position	Object name	Value
1	holger2	19.55 MB
2	holger9	17.57 MB
3	john9	15.00 MB
4	john11	11.43 MB
5	holger1	7.54 MB
6	holger3	6.21 MB
7	john12	4.80 MB
8	john13	4.48 MB

Current Status - SMBTAmonitor

- Connects directly to SMBTAd, to retrieve real-time information on the data flow
- ncurses based




The screenshot shows a terminal window titled "holger@dl33.suse.de: ~/t/smbttools/run". The application is "SMBTAmonitor 1.2.3" monitoring the object "bss-data". It displays a summary of data flow and a log of file operations.

Total (Read/Write)	Total (Reading)	Total (Writing)
929.84 MB (975010209 B)	360.21 MB (377708463 B)	220.41 MB (231112283 B)
3.76 KB / second	3.76 KB / second	0.00 Bytes / second


```
2011-03-04 13:24:31.876: User john opened file shared/persons/flat_data/ircbot.py on share bss-data (LINUX-MAMY).
2011-03-04 13:24:30.896: User holger closed file datapool/special/classic/index.html on share bss-data (LINUX-MAMY).
2011-03-04 13:24:30.895: User holger read from file datapool/special/classic/index.html on share bss-data (LINUX-MAMY).
2011-03-04 13:24:30.894: User brtam closed file homepage/all/tomboy-notes-backup-180909.tar.bz2 on share bss-data (LINUX-MAMY).
2011-03-04 13:24:30.894: User holger read from file datapool/special/classic/index.html on share bss-data (LINUX-MAMY).
```

Current Status - rrddriver

- Rrdtool is a well known round-robin-database to create usage graphs and statistics
- Rrddriver forms an interface from SMBTAd to rrdtool
- Connects to the real-time interface like SMBTAmomitor.



Current Status - webSMBTA

- Our latest development, a web Interface for SMBTAquery
- Rails3 based
- Developed by Benjamin Brunner as his final exam for his apprenticeship at SUSE

The screenshot displays the webSMBTA interface, version 1.2.3. It features a green header with the application logo and name. Below the header, there are links for 'Homepage' and 'Online help'. The main content area is divided into four panels, each representing a different level of the SMBTA hierarchy: 'Global', 'Domains', 'Shares', and 'Files'. The 'Global' panel shows a list with 'global' selected. The 'Domains' panel shows '(All)' and 'SEIZEWELL'. The 'Shares' panel shows '(All)', 'tracking', and 'trigger'. The 'Files' panel shows a long list of files and folders, including 'datapool/assigned/dvd_copy_finals_2012.iso', 'datapool/single/tomboy-notes-backup-180909.tar.bz2', 'homespace/all/battery_charging_time.jpg', 'office_space/data/y2010/call-minutes-17042011.txt', 'office_space/transfer/bss/guide-oem-2-11.pdf', 'shared/appdata/celldriven.doc', 'shared/appdata/guide-reseller.pdf', 'shared/appdata/samba-3.0.32.tar.bz2', 'shared/appdata/setup.exe', 'shared/global/save-2.sav', 'shared/logged/pict0001.png', 'shared/programs/excelsior-contract-2011.doc', 'shared/programs/opensuse-tumbleweed.iso', 'src1/work/needed/ebermannstadt-index.html', and 'src1/work/needed/log.smbd'. A 'Function:' dropdown menu is located below the 'Global' panel, with a prompt: 'Choose an object from the tables above, and select a function'. Blue arrows indicate the flow from one panel to the next.

Current Status - webSMBTA

- Our latest development, a web Interface for SMBTAquery
- Rails3 based
- Developed by Benjamin Brunner as his final exam for his apprenticeship at SUSE



Project data and outlook



Project outlook

- Project has reached stable grounds:
 - SMBTAd is used in production on some sites
 - We transferred tens of Terabytes through SMBTAd without getting it to crash in testsuites
- Implement SMBTAmonitor in GTK/QT/Java
 - Integrate with webSMBTA
- Support for clustered Samba inside the module
 - Be able to watch how a cluster distributes Samba traffic on the nodes

Project data, Q&A

- Homepage
 - <http://holger123.wordpress.com/smb-traffic-analyzer/>
- Defect and version tracking:
 - <http://bugzilla.samba.org>
 - > Product “smbta”
- Discussion
 - Mailinglists: samba@samba.org, or samba-technical@samba.org
 - IRC: [irc.freenode.net](irc://irc.freenode.net), channel [#smbta](#), [#samba](#), and [#samba-technical](#)
- SMBTA Team:
 - Holger Hetterich (I have my fingers in any of the stuff)
 - Benjamin Brunner (is developing webSMBTA and SMBTAquery)
 - Michael Häfner (works on rrd driver and SMBTAmonitor)
 - Robert Piasek (extensive Testing, occasional bugfixing, Gentoo Packager)
 - Björn Geuken (developing a WebYaST module to control SMBTAd)

