A year with the Microsoft Docs

The Good, the Bad and the Ugly
of Samba4 and Microsoft
SambaXP 2009

Andrew Bartlett <abartlet@samba.org>
Samba Team
Red Hat Inc.

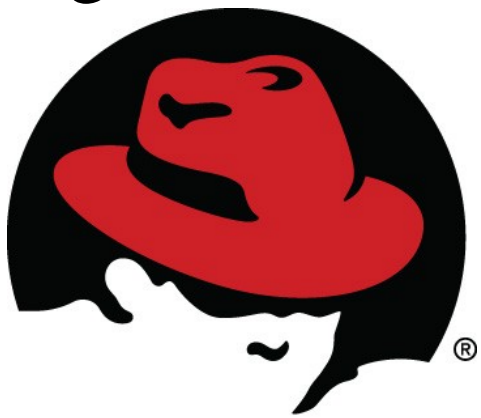

abartlet

- A Samba developer for 8 years now
- One of the lead developers on Samba4
- I work for Red Hat
  - Full time on Samba4
  - But these are my views, not theirs
- **Please** ask questions during the talk

Microsoft

- Our partners in this dance
- Without them Samba would not be where it is
  - Is that a good thing?
- Now provides protocol documentation
  - As required by the European Courts
  - But also more than as required – like in public
  - I work closely with them on our doc issues/questions

Samba4

- A replacement for Active Directory
  - Provides AD-compatible such as a:
    - KDC
    - Netlogon server
    - LDAP server
    - LSA server
    - SAMR server
- My primary work area
  - So my questions and interactions are AD related

The Good

Documentation from Microsoft

- At SambaXP last year we had Microsoft's docs in private
  - Special agreements had to be signed
  - I am a PFIF subcontractor
    - This gave me the right to see them under NDA
- Now (May 2008) everyone has them
  - **Not** required by the EU decision
  - Now over 300 protocols publicly documented

NETLOGON Datagrams

- My first real use of the WSPP docs
- Packet used to locate a domain controller
  - Encoded over CLDAP and NetBIOS
- Previously had 3 different parsers
  - Docs make the links and commonality clear
  - Now just a single parser, even shared with Samba3
- Microsoft fixed issues I found from the implementation

Microsoft-compatible NTP

- Windows clients need time synchronisation
- Microsoft made a custom extension to NTP
  - Unrelated to the standard NTP authentication
  - Even used the opposite bit order!
  - Needs a custom patch to NTP
- Without the docs, we would be back to crypto guessing games
  - They were fun, but not productive

AD/LDAP semantics

- Which attributes are generated?
- How are they generated?
  - possibleSuperiors in particular
- How does the schema work?

Plugfests!

- Microsoft put on two plugfests in Sep/Oct 2008
  - Sponsored the SNIA plugfest in San Jose for the CIFS industry
    - Particularly special because for many years, Microsoft would not even attend
  - Hosted a special 'Samba Team' event for us at Redmond


Microsoft

A week working with Microsoft

- Knowledgeable testers, running Microsoft's own tests
  - Found many, many bugs
- Original developers, chasing down harder bugs
  - Developers were willing to spend all day debugging some of our trickier issues.

More than documentation

- We have real help from Microsoft!
  - I disagreed with Microsoft on parsing a blob
  - So Richard Guthrie hand-parsed it!

01 00 00 00 0C 00 00 00 38 00 00 00 F8 03 08 26

7A 22 C9 01 02 00 00 00 1C 00 00 00 89 F7 2F 75

14 99 A1 70 38 73 DA ED 55 D3 1E EE 78 82 B7 A8

6F 6C 23 E6 83 83 00 1F 01 00 00 00

Count of outgoing auth infos (AuthInfos) = 1

Byte offset to outgoing AuthenticationInformation = 12

Byte offset to outgoing PreviousAuthenticationInformation = 56

LSAPR_AUTH_INFORMATION – auth outgoing current authentication info

LastUpdateTime = F80308267A22C901

AuthType = 2

AuthInfoLength = 24

AuthInfo = 89F72F751499A1703873DAED55D31EEE7882B7A86F6C23E68383001F

Count of incoming auth infos (AuthInfos)= 1

The Bad

Some bugs take time

- Windows tries to change it's password monthly
  - It does so by filling a unicode (UTF16) buffer with random data
  - Then sending it (encrypted) to the server
- How do you convert random data to UTF8?
  - Bad things happen when you convert it to “” instead
  - Required a rework of the whole password-setting stack
  - Special handling for invalid sequences

Some things take a lot of time

- It was over 6 months to get a correct text-file of the AD schema
  - Not all of it was waiting on Micorosoft
  - But Microsoft's own developers don't have to do this
- Shows how much this is is still 'document later'

And some just want a little validation

- My 'russian connection' had a strange error
  - 'PAC Validation failed'.
- PAC: Privilage Attribute Certificate
  - Microsoft's extension to attach groups to a Kerberos ticket
  - Windows XP must check the PAC with the KDC
  - But only rarely
    - so I never saw it in my testing
- This is why real-world testing is vital

Challenges in the Docs

- The docs are certainly not perfect
- Many produced by an 'archaeological' process
  - The original authors have long left the company
  - There were no similar docs created at the time
- Using the documentation takes more time
  - Much harder to do a half-job when the whole task is clear
- It takes time to peruse the corrections and clarifications

The Ugly

How hard to get a text file?

- It took 6 months to get a correct copy of the AD schema file
  - Not all Microsoft's delay, but still unacceptable
  - Why could we not just get the same file as imported into windows?
- We never expected a level playing field
  - But this shows the tilt very well

The Lawyers are watching!

- I've had a fair bit to say about the whole docs process
    - I work regularly with a team of enthusiastic engineers who clearly want us to succeed ... these engineers do everything they can to provide us with the information we need.
    - However, the Microsoft documents have been written so as to frustrate the implementer at every turn.
- Both quotes are now in evidence in the US courts!

Patent Sabar Ratting continues

- We see from the TomTom case that Microsoft **will** go after Linux
- We also know that most patent threats and deals are never public
- We must remain vigilant
  - Our great relationship with the Protocol Program team won't save us from the rest of the company
  - We do have some protections in the WSPP agreement
    - This is a binding list of applicable patents

Parsing bitfields

- Who can tell me what the value of bit Q is?
  - From 3.5.4.2.1 DsrGetDcNameEx2

Flags: A set of bit flags in little-endian format that provide additional data that is used to process the request. A flag is TRUE (or set) if its value is equal to 1. The value is constructed from zero or more bit flags from the following table, with the exceptions that bits D, E, and H cannot be combined; S and R cannot be combined; and N and O cannot be combined.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
S	R	0	0	0	0	0	0	0	0	0	T	Q	P	O	N	M	L	K	J	I	H	G	F	E	D	C	B	0	0	0	A

A: Forces cached DC data to be ignored.

B: Requires that the returned DC support specific operating system versions. <158>

Q: Requires that the returned DC be running a specific operating system. <162>

Conclusion:
How can this improve?

Fix the formats:

```
typedef [bitmap32bit] bitmap {
    DS_FORCE_REDISCOVERY       = 0x00000001,
    DS_DIRECTORY_SERVICE_REQUIRED = 0x00000010,
    DS_DIRECTORY_SERVICE_PREFERRED   = 0x00000020,
    DS_GC_SERVER_REQUIRED      = 0x00000040,
    DS_PDC_REQUIRED          = 0x00000080,
    DS_BACKGROUND_ONLY       = 0x00000100,
    DS_IP_REQUIRED           = 0x00000200,
    DS_KDC_REQUIRED          = 0x00000400,
    DS_TIMESERV_REQUIRED       = 0x00000800,
    DS_WRITABLE_REQUIRED       = 0x00001000,
    DS_GOOD_TIMESERV_PREFERRED= 0x00002000,
    DS_AVOID_SELF        = 0x00004000,
    DS_ONLY_LDAP_NEEDED      = 0x00008000,
    DS_IS_FLAT_NAME          = 0x00010000,
    DS_IS_DNS_NAME           = 0x00020000,
    DS_TRY_NEXTCLOSEST_SITE        = 0x00040000,
    DS_DIRECTORY_SERVICE_6_REQUIRED = 0x00080000,
    DS_RETURN_DNS_NAME       = 0x40000000,
    DS_RETURN_FLAT_NAME      = 0x80000000
} netr_DsRGetDCName_flags;
```

This would be better too:

Q: Requires that the returned DC be running a specific operating system. <162>

0x0008 0000 DS_DIRECTORY_SERVICE_6_REQUIRED

Eliminate the humans

- The schema and IDL documents should be automatically generated
  - There should be no room for human 'corrections'
    - There are 'errors' in the IDL that must be preserved for compatibility
  - Supply us the same source format used by MS

So are we more productive?

- We are more productive with the Microsoft docs
- But we must still prove the documentation
  - To use the docs without a testsuite would be foolish
  - Sometimes they are incorrect
- There has not been a sudden rush of new developers

More eyeballs

- We need more help checking the docs
- Also help comparing Samba with the docs
- You should not have to be a Samba Wizard
  - They docs should be understandable by mortals too
  - If you don't understand it, Microsoft should clarify it

Thanks / Q&A

- Are there any (more) questions?