



IBM

Testing CTDB

- not necessarily trivial!

- **Martin Schwenke** • martins@au1.ibm.com • martin@meltin.net
- **Ronnie Sahlberg** • sahlberg@au1.ibm.com • ronniesahlberg@gmail.com
-
- **IBM** • Australia Development Laboratory • Linux Technology Center

18 January 2012

IBM



The Plan

- What is CTDB?
- What needs to be tested?
- The old “test suite”
- Testing eventscripts using stub command scripts
- CCAN-style tests
 - Just include everything!
 - ... versus (shared) library
 - For testing parts of ctdbd and ctdb tool
- System testing using Autocluster
- Performance testing with DBENCH



What is CTDB?

- ctdb.samba.org
- Clustered Trivial Database (based on TDB)
- All-active multi-node cluster
- Functions:
 - Database
 - Monitoring
 - IP Address Failover
- For more details see Tridge's LCA2007 talk



What needs to be tested?

- Everything!
- This is complex, so we need regression tests...
- Test the simple things first!
- Unit testing makes failures easy to diagnose
- Complex system/cluster tests last...
- Performance regressions
- Memory leaks
- File descriptor leaks



The old “test suite”...

- Tests the basic functions of the CTDB command-line tool in combination with ctddb
- Quite a few tests removed over time - they were stupid!
- A few failover tests
- This test suite is really good at finding races and other bugs in the startup code, since it restarts ctddb a lot



Eventscripts - system stubs

- Eventscripts are shell scripts that monitor and modify the system for ctdbd
- Eventscripts use a lot of system commands
 - ip, iptables, ethtool, sleep, netstat, kill(all), rpcinfo, exportfs,
- Eventscripts use ctdb command-line tool
- Rewrote huge slabs of the eventscripts, had no regression tests... how to test in isolation?
- Write scripts that fake behaviour and/or output for the parts of these commands that are interesting and/or required
- Examples...



IP allocation algorithm

- CTDB assigns one or more groups of IP addresses across nodes in a cluster
- Wanted to add a new algorithm so wrote prototype in Python
- Ran within Python simulation, looked good...
- Rewrote in C for ctddb
- How to unit test?
- Don't make the functions static?
- We'll need more of these test programs...
shared library?



CCAN style tests

- CCAN style testing: just #include everything!
- Static functions available :-)
- Static functions clash :-(
- Not that hard to resolve...
- Easy to write a testcase for the first bug found
- Examples...
- Python simulation now calls binary with ctddb's IP allocation algorithm...



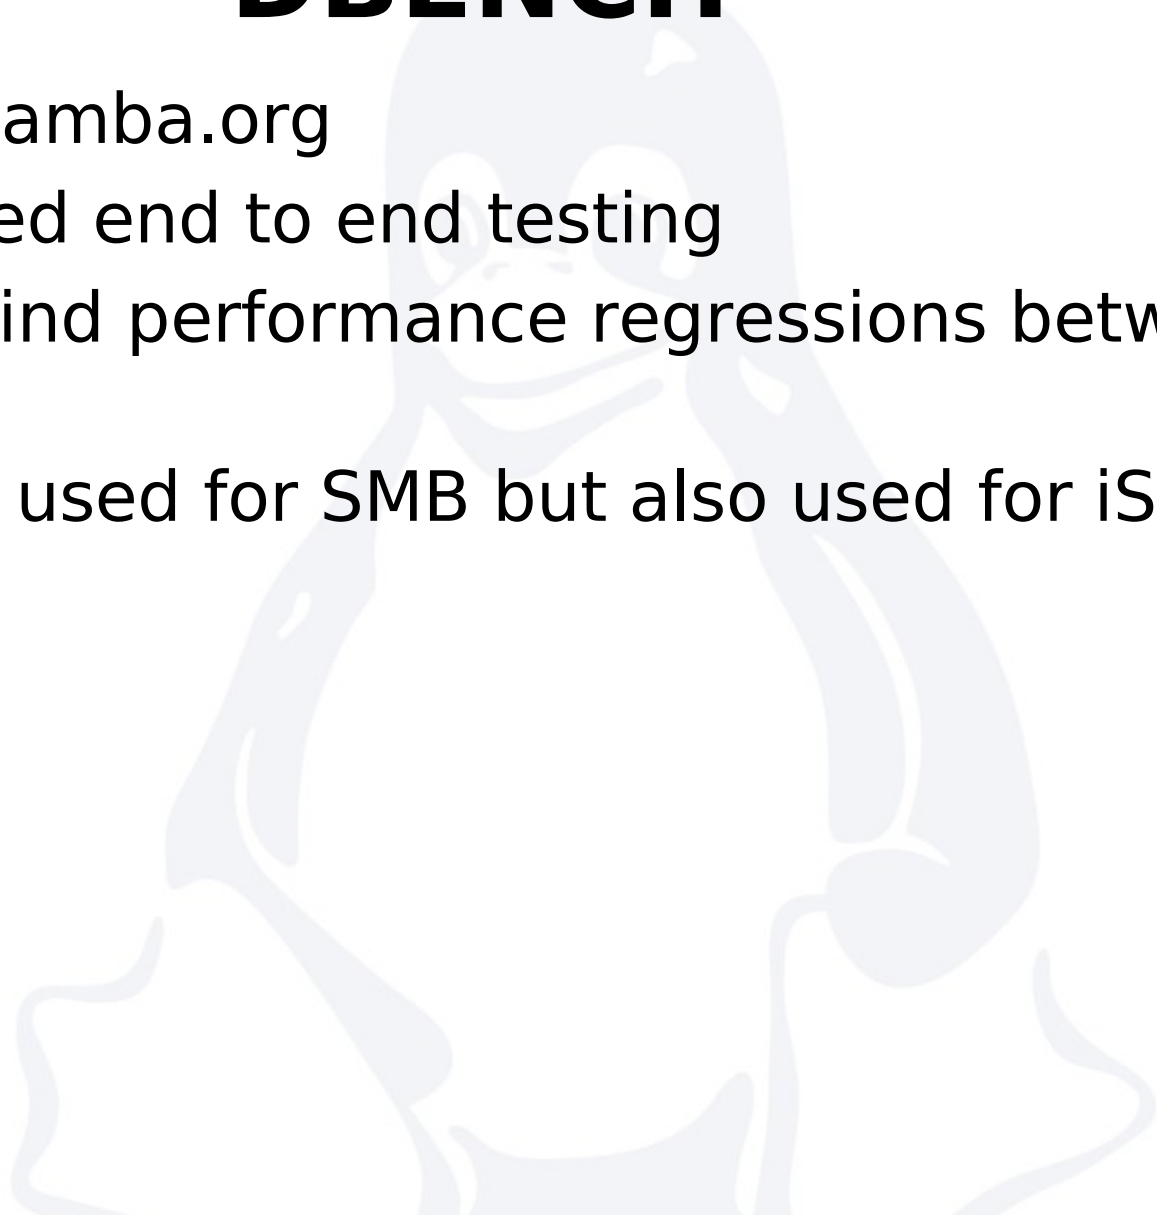
Autocluster

- Autocluster is a tool for testing clustered Samba
- Autocluster is a huge, configurable set of shell scripts for creating and configuring virtual KVM-based clusters
- See LCA2009 talk with Tridge
- Added new template-based service configuration subsystem - canned configuration kept on bit-rotting
- Used for developer testing of CTDB, clustered Samba, clustered NFS, ...
- Used as basis for product regression testing



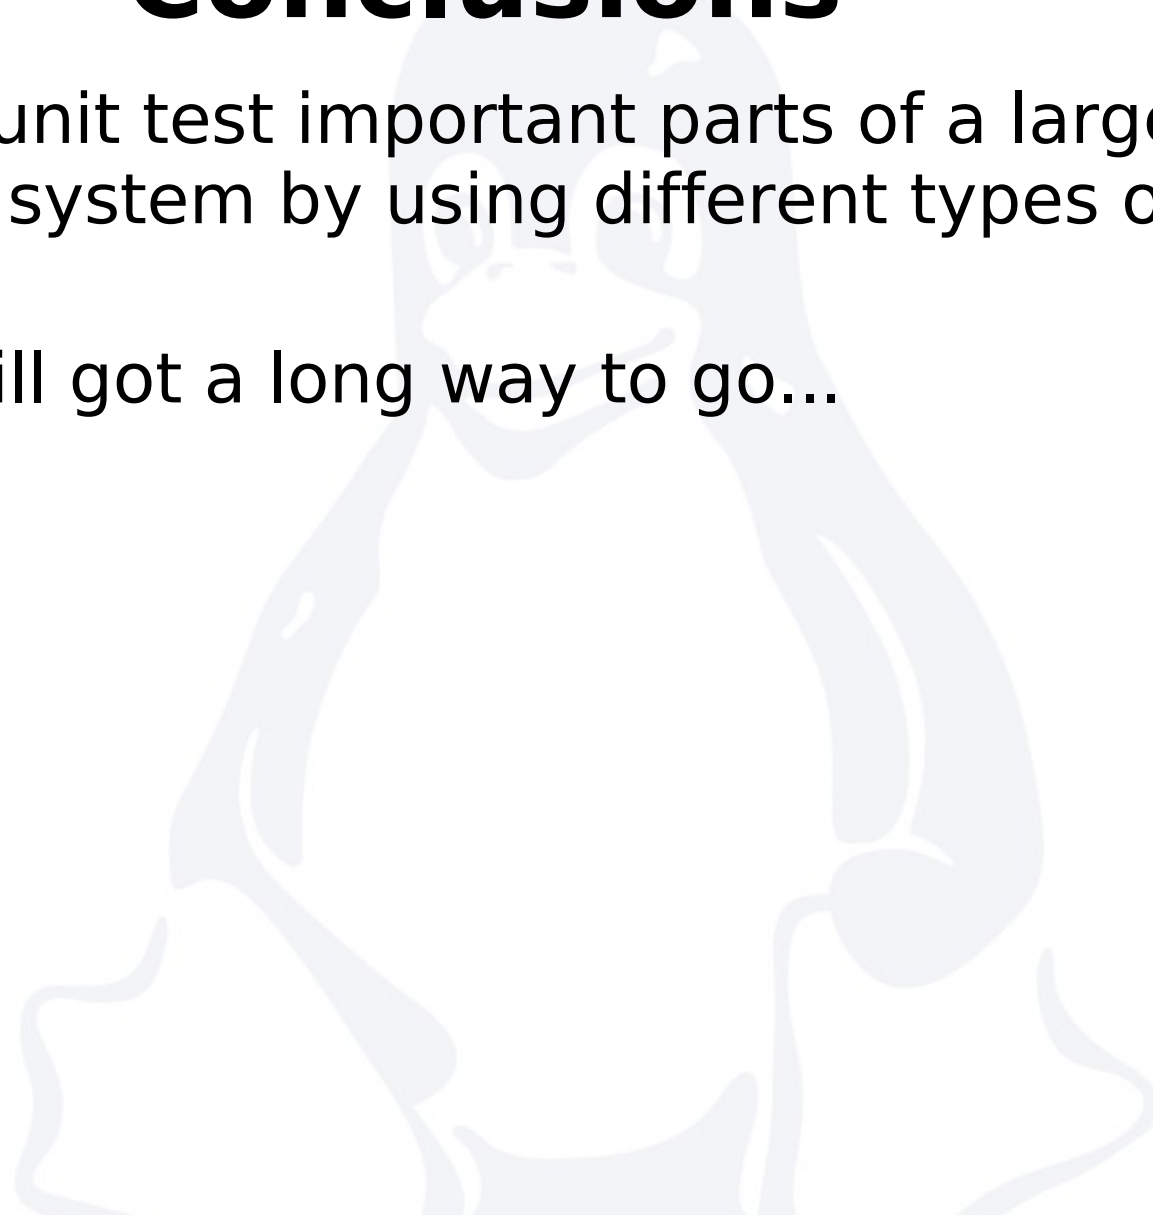
DBENCH

- dbench.samba.org
- Automated end to end testing
- Used to find performance regressions between versions.
- Primarily used for SMB but also used for iSCSI, and NFS.



Conclusions

- You can unit test important parts of a large, complex system by using different types of stubs
- We've still got a long way to go...



Q & A

Also feel free to grab us later if you want to chat...

IBM



Legal Statement

- This work represents the view of the authors and does not necessarily represent the view of IBM.
- IBM is a registered trademark of International Business Machines Corporation in the United States and/or other countries.
- Linux is a registered trademark of Linus Torvalds.
- Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Other company, product, and service names may be trademarks or service marks of others.

