



# The Aspersa Toolkit

Baron Schwartz  
O'Reilly MySQL Conference & Expo 2011



**Consulting  
Support  
Training  
Development**

**For MySQL**

# Percona Server



- Replaces MySQL
- Faster Queries
- More Consistent
- More Measurable
- More Features

# Percona XtraBackup



PERCONA  
XTRABACKUP

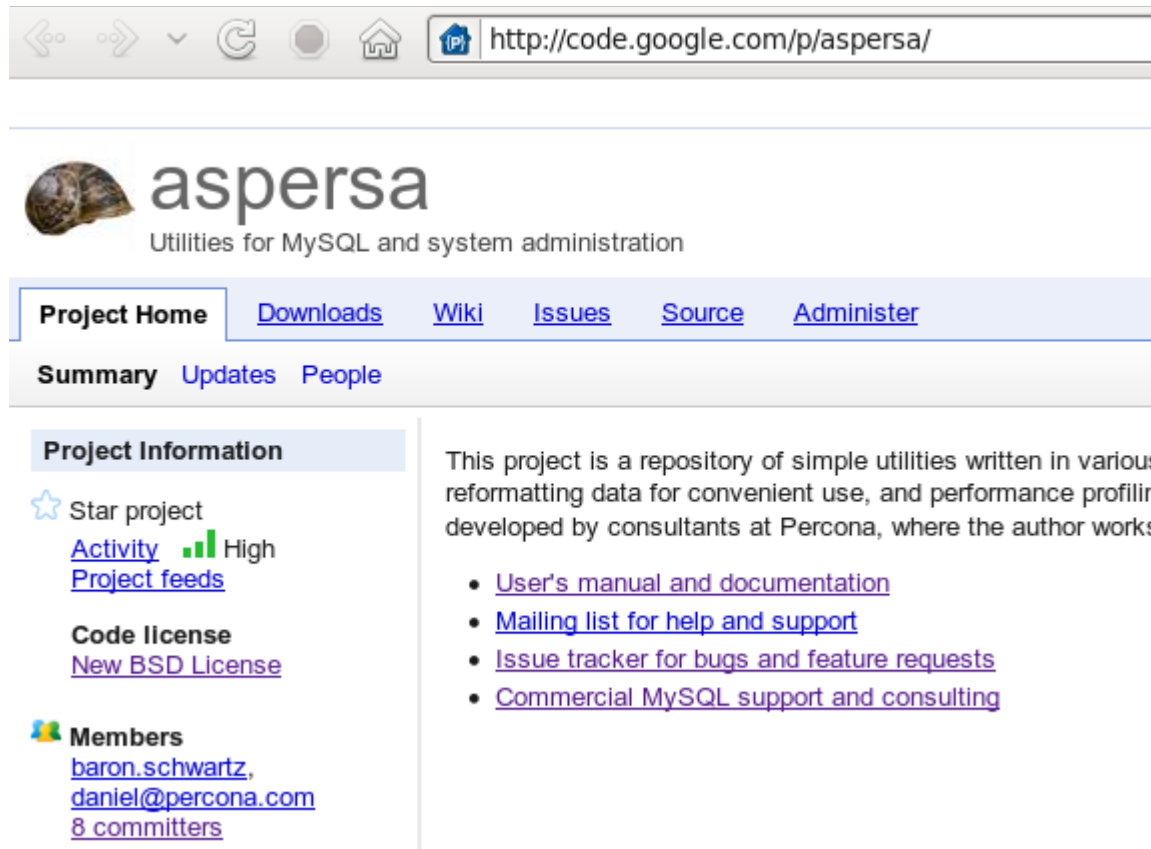
- Backs Up InnoDB
- Non-Blocking

# What's Aspersa?

- It's a snail, and the target of Twitter spam
- No, seriously, it's Bash (and some Perl) tools

# Aspersa On The Web

- <http://code.google.com/p/aspersa>




The screenshot shows a web browser window with the address bar containing <http://code.google.com/p/aspersa/>. The page features the Aspersa logo, which is a snail, and the text "aspersa Utilities for MySQL and system administration". Below the logo is a navigation menu with links for "Project Home", "Downloads", "Wiki", "Issues", "Source", and "Administer". Under "Project Home", there are sub-links for "Summary", "Updates", and "People". The "Project Information" section includes a "Star project" button, an "Activity" link with a green bar chart and the word "High", and a "Project feeds" link. The "Code license" section lists "New BSD License". The "Members" section lists "baron.schwartz", "daniel@percona.com", and "8 committers". The main content area contains a paragraph describing the project as a repository of simple utilities for MySQL and system administration, developed by consultants at Percona. Below this paragraph is a bulleted list of links: "User's manual and documentation", "Mailing list for help and support", "Issue tracker for bugs and feature requests", and "Commercial MySQL support and consulting".


Project Home Downloads Wiki Issues Source Administer

Summary Updates People

**Project Information**

★ Star project  
[Activity](#)  High  
[Project feeds](#)

**Code license**  
[New BSD License](#)

 **Members**  
[baron.schwartz](#),  
[daniel@percona.com](#)  
8 committers

This project is a repository of simple utilities written in various reformatting data for convenient use, and performance profiling developed by consultants at Percona, where the author works.

- [User's manual and documentation](#)
- [Mailing list for help and support](#)
- [Issue tracker for bugs and feature requests](#)
- [Commercial MySQL support and consulting](#)

# Today's Agenda

- diskstats
- ioprofile
- summary
- mysql-summary
- stalk, collect, and sift

# Installing Aspersa Tools

- Don't install: just wget, or curl, or fetch, or GET.
- `wget aspersa.googlecode.com/svn/trunk/<tool>`

# The diskstats tool

- An iostat replacement for Linux
- Better than iostat!

#ts	device	rd_s	rd_avkb	rd_mb_s	rd_mrg	rd_cnc	rd_rt	wr_s	wr_avkb	wr_mb_s	wr_mrg	wr_cnc	wr_rt	busy	in_prg
1.0	{28}	74.4	41.3	1.5	0%	0.3	94.5	62.5	14.1	0.4	43%	0.0	4.6	1%	0
2.0	{28}	168.2	59.4	4.9	0%	0.1	23.2	92.0	14.9	0.7	46%	0.0	1.8	3%	0
3.0	{28}	218.3	56.1	6.0	0%	0.3	38.4	418.8	15.3	3.1	48%	0.0	2.5	6%	0
4.0	{28}	79.3	52.8	2.0	0%	0.1	24.9	153.5	27.8	2.1	71%	0.1	9.5	2%	0
5.0	{28}	229.8	61.4	6.9	0%	0.3	35.0	70.3	13.4	0.5	40%	0.0	1.3	3%	0
6.1	{28}	226.0	62.0	6.8	0%	0.3	34.7	65.4	12.6	0.4	37%	0.0	5.3	3%	0
7.1	{28}	136.9	61.9	4.1	0%	0.1	22.2	9.9	8.0	0.0	0%	0.0	27.0	2%	6
8.1	{28}	155.8	55.0	4.2	0%	0.1	18.9	49.6	13.1	0.3	39%	0.0	1.1	2%	9
9.1	{28}	52.6	42.4	1.1	0%	0.0	24.7	126.0	18.9	1.2	58%	0.0	4.4	2%	58
10.1	{28}	311.4	60.7	9.2	0%	0.4	35.1	83.3	14.5	0.6	45%	0.0	1.0	6%	0
11.1	{28}	159.3	60.5	4.7	0%	0.2	34.3	99.3	14.7	0.7	45%	0.0	1.6	3%	0
12.1	{28}	176.5	58.7	5.1	0%	0.2	30.4	99.1	10.9	0.5	28%	0.0	0.8	3%	2
13.1	{28}	246.0	57.9	7.0	0%	0.3	37.0	108.1	13.4	0.7	40%	0.0	1.4	5%	0
14.1	{28}	176.5	61.8	5.3	0%	0.3	45.3	88.2	19.5	0.8	68%	0.0	9.4	4%	12
15.1	{28}	35.7	20.4	0.4	0%	0.0	5.9	2809.5	16.7	22.9	52%	0.5	5.2	6%	0
16.1	{28}	101.2	58.5	2.9	0%	0.0	10.8	75.4	14.1	0.5	43%	0.0	2.4	1%	0
17.2	{28}	91.1	59.1	2.6	0%	0.0	12.2	95.1	11.1	0.5	28%	0.0	3.4	1%	0
18.2	{28}	99.2	55.0	2.7	0%	0.1	16.4	73.4	12.3	0.4	35%	0.0	5.5	2%	0
19.2	{28}	91.2	56.7	2.5	0%	0.1	26.3	111.0	20.6	1.1	61%	0.0	7.6	2%	0
20.2	{28}	94.9	57.0	2.6	0%	0.0	11.4	61.3	12.6	0.4	37%	0.0	2.3	1%	0
21.2	{28}	81.3	57.2	2.3	0%	0.1	26.0	57.5	14.6	0.4	45%	0.0	1.1	2%	0
22.2	{28}	228.7	59.9	6.7	0%	0.2	27.7	17.8	8.0	0.1	0%	0.0	5.8	4%	19
23.2	{28}	36.7	36.1	0.6	0%	0.1	100.2	65.4	13.8	0.4	42%	0.0	4.7	0%	0
24.2	{28}	95.1	55.2	2.6	0%	0.0	12.2	109.9	19.7	1.1	59%	0.0	5.9	1%	0
25.2	{28}	93.2	58.6	2.7	0%	0.0	12.1	68.4	13.8	0.5	42%	0.0	2.2	1%	0
26.2	{28}	125.0	48.8	3.0	0%	0.1	22.3	2810.6	16.4	22.5	51%	0.6	6.2	8%	0
27.2	{28}	25.8	25.2	0.3	0%	0.0	1.8	84.3	14.0	0.6	43%	0.0	1.3	0%	0
28.3	{28}	270.6	60.6	8.0	0%	0.1	11.7	76.3	14.4	0.5	45%	0.0	1.8	3%	2
29.3	{28}	23.7	17.0	0.2	0%	0.0	9.2	113.8	20.0	1.1	60%	0.0	7.1	1%	1

# Live diskstats demo

# Diskstats If You Missed The Demo

- Interactive, key-driven operation
- Separate stats for reads and writes
- Correct utilization and busy-ness statistics
- Easy to drill down into disks and time periods
- Trivial to collect data for later usage
- Lightweight, no need to “install”
- No dependencies – just awk/grep/etc
- I prefer it even when iostat is installed!

# The ioprofile Tool

- Uses strace and lsof to see what files consume a process's IO time.
- Be careful with production servers; strace has been known to leave them in a weird state.

# Live ioprofile demo

# lprofile If You Missed The Demo

- Prints out a profile of files and function calls
- MySQL 5.5 has the Performance Schema, which is the Right Way To Do It instead.
- Simplified screenshot:

```
total pwrite64      fsync filename
0.328388 0.001143 0.327245 /home/baron/etc/mysql/server/5.1.50/data/ib_logfile0
0.094027 0.000082 0.000000 /home/baron/etc/mysql/server/5.1.50/data/test/t.frm
0.086538 0.000970 0.085568 /home/baron/etc/mysql/server/5.1.50/data/ibdata1
0.014091 0.000000 0.000000 /home/baron/etc/mysql/server/5.1.50/data/test/db.opt
0.009051 0.000000 0.000000 /home/baron/etc/mysql/server/5.1.50/data/test/
0.000602 0.000000 0.000000 /home/baron/etc/mysql/server/5.1.50/data/ginger-bin.000011
0.000463 0.000000 0.000000 /home/baron/etc/mysql/server/5.1.50/data/
0.000085 0.000000 0.000000 /home/baron/etc/mysql/server/5.1.50/data/mysql/db.opt
```

# The summary Tool

- Pretty-formats a system overview
- Easy to email
- Easy to diff
- Does lots of things like checking RAID status

Live “summary” demo

# The mysql-summary Tool

- Pretty-prints a MySQL server overview
- Similar goals to “summary” (diffable, etc)

# Live “mysql-summary” demo

# The stalk/collect/sift triad

- Capture system information when triggered
- Extremely helpful for diagnosing problems
- Very MySQL-oriented
- Three parts:
  - Stalk looks for something “wrong” and fires collect
  - Collect captures tons of information to files
  - Sift lets you browse these files easily

# Why So Important?

- Diagnosing problems is trivial when:
  - You have all the data you need.
  - It is properly scoped.
  - You have good tools for looking at it.
  - You know what you're looking at.
- When you don't have good data?
  - Simple problems drag on for months.

# Don't Roll Your Own

- This tool combination sounds simple, but isn't
- Lots of little details are important to get right
  - Example: don't collect data if it'll fill up the disk

# The Most Important Step

- Correctly defining “abnormal” is vital
- If you don't get this right, you get
  - False positives: waste of time, creates doubt
  - False negatives: you don't resolve the problem
- The tool *must* be configured
  - There is no one-size-fits-all configuration!

# How stalk works

- The stalk tool runs a check periodically
- If the results are out of range, it fires collect
- The default check is `Threads_connected > 100`
  - This must be customized!
- Create a `stalk.conf` file for configuration

# Starting Stalk

- Open a “screen” session
- Run “stalk” as superuser
- Wait for results. (You can make it email you.)

# Interpreting Results

- Collect collects a LOT of data.
- Wading through it by hand is miserable.
- Use the sift tool.
  - Start it with the directory where collect saved data.
  - Use interactive key-driven functionality to browse.
  - Zoom in on places that look troublesome.

# Live “sift” demo

# Sift, If You Missed The Demo

- Use “vim” keys k and j to navigate between samples
- Use the ? key to find out what keys are active
- Dig into samples that show something wrong
  - If you set up your config right, this is hopefully the first sample :)

# Stalk/Collect Long-Term

- This started as an install-when-needed tool
- It might be a good idea to let it run ALL the time
  - It auto-purges old samples and won't fill your disk
- Example:
  - My server's Threads\_running is always 5 to 15
  - Set it up to capture when Threads\_running > 25

# Percona Live, May 26, New York



[www.percona.com/live](http://www.percona.com/live)



**baron@percona.com**

**We're Hiring! [www.percona.com/about-us/careers/](http://www.percona.com/about-us/careers/)**