

# **Claims Of Corporate And Legal Misconduct**

## **Part I Of II — Acts Of Fritz Knabe**



Walter Tuvell

August 18, 2011

## Document History

Rev.	Date	Author	Remarks
0.1	August 4, 2011	Walter Tuvell	Emergency draft, in case Dan Feldman, <i>et al.</i> , succeed in terminating me before final 1.0, preventing full C&A Process
1.0	August 18, 2011	Walter Tuvell	Split into Parts I & II; Final 1.0 version

## Related Documents

Author	Date	Title
IBM	2011	<i>Business Conduct Guidelines (BCG)</i>
IBM	June, 2009	<i>About Your Job (AYJ)</i> , Document No. USHR101
IBM	May 18, 2008	<i>Concerns and Appeals Program (C&amp;A)</i> , Document No. USHR102
Cynthia Shapiro	2005	<i>Corporate Confidential: 50 Secrets Your Company Doesn't Want You to Know — and What to Do About Them</i> , St. Martin's Griffin (ISBN-13: 978-0312337360)
Michelle Crouch	April, 2011	<i>Get Hired, Not Fired — 50 Secrets Your HR Person Won't Tell You</i> , cover story of Reader's Digest, pp. 130-139; with an online supplement at <a href="http://www.rd.com/13-things/39-more-secrets-your-hr-person-wont-tell-you/">http://www.rd.com/13-things/39-more-secrets-your-hr-person-wont-tell-you/</a>
Wikipedia	February, 2006 – Present	<i>Workplace bullying</i> , online at <a href="http://en.wikipedia.org/wiki/Workplace_bullying">http://en.wikipedia.org/wiki/Workplace_bullying</a>
Kevin Maney, Steve Hamm, Jeffrey M. O'Brien,	2011	<i>Making the World Work Better, The Ideas That Shaped a Century and a Company (MTWWB)</i> , IBM Press — Pearson plc

## Table of Contents

Document History .....	2
Related Documents .....	2
1 Executive Summary — Part I .....	5
1.1 Disclaimer .....	5
1.2 Theory Of The Case: Blackballing .....	6
1.3 List Of Particulars .....	8
2 Background — Part I .....	10
2.1 Introduction .....	10
2.2 Relations With Dan And Fritz .....	11
2.3 Wahoo Instability .....	11
2.4 Other Work .....	13
3 Fritz: Database Transport (March 15-17) .....	14
3.1 Aftermath .....	15
4 Fritz: Excel Graphics (May 18) .....	15
4.1 Aftermath .....	18
5 Fritz: Wahoo Not The Bottleneck (June 6) .....	20
6 Fritz: Yelling In Public (June 8) .....	21
6.1 Aftermath .....	24
APPENDICES — Part I .....	27
A Weekly Reports .....	27
A.a Week Of Nov 8-14 (2010) .....	27
A.b Weeks Of Nov 17-21, Nov 24-28 (2010) .....	28
A.c Week Of Nov 29 - Dec 3 (2010) .....	28
A.d Week Of Dec 6-12 (2010) .....	28
A.e Week Of Dec 13-19 (2010) .....	29
A.f Week Of Dec 20-26 (2010) .....	30
A.g Week Of Dec 27, 2010 - Jan 2, 2011 .....	30
A.h Week Of Jan 2-9 .....	31
A.i Week Of Jan 9-16 .....	32
A.j Week Of Jan 16-23 .....	32
A.k Week Of Jan 23-30 .....	32
A.l Week Of Jan 30 - Feb 6 .....	33
A.m Week Of Feb 6-13 .....	33
A.n Week Of Feb 13-20 .....	34
A.o Week Of Feb 20-27 .....	35
A.p Week Of Feb 27 - Mar 6 .....	35
A.q Week Of Mar 6-13 .....	36
A.r Week Of Mar 13-20 .....	36
A.s Week Of Mar 20-27 .....	37

A.t	Week Of Mar 27 - Apr 3 .....	37
A.u	Week Of Apr 3-10 .....	38
A.v	Week Of Apr 10-17 .....	38
A.w	Week Of Apr 17-24 .....	39
A.x	Week Of Apr 24 - May 1 .....	39
A.y	Week Of May 1-8 .....	40
A.z	Week Of May 8-15 .....	40
A.aa	Week Of May 16-22 .....	40
A.bb	Week Of May 22-29 .....	41
A.cc	Week Of May 29 - Jun 5 .....	41
A.dd	Week Of Jun 5-12 .....	42
A.ee	Week Of Jun 12-19 .....	43
A.ff	Week Of Jun 19-26 .....	44
A.gg	Week Of Jun 26 - Jul 3 .....	44
A.hh	Week Of Jul 3-10 .....	45
A.ii	Week Of Jul 10-17 .....	46
A.jj	Week Of Jul 17-24 .....	46
A.kk	Week Of Jul 24-31 .....	46
A.ll	Week Of Jul 31 - Aug 7 .....	46
A.mm	Week Of Aug 7-14 .....	46
B	Email Chain: Database Transport (March 15-16) .....	46
C	Email Chain: Excel Graphics (May 19) .....	51
D	Email Chain: Wahoo Not The Bottleneck (May 31) .....	53
E	Email Ensemble: Continuing Untestability Of Wahoo (May 23 - June 10) .....	54
F	Email Chain: Next Steps (June 2) .....	108
G	Report And Email Chain: Wahoo 4× Faster Than Skimmer (June 8-9) .....	111
G.a	PerfScore Report (June 8) .....	112
G.b	Follow-Up Email Chain (June 8-10) .....	114
H	Email Chain: Upcoming Performance Tests (June 8-10) .....	115
I	Email Chain: Yelling Incident (June 8-10) .....	117
I.a	Comments On Some “Miscommunications” (June 8) .....	117
I.b	Follow-Ups And Apologies (June 10) .....	119
J	Email Chain: Request For Help (June 10) .....	120
K	Sujatha’s Work Items .....	124
L	Email Chain: Hail And Farewell (June 11-13) .....	125
L.a	Goodbye Letter .....	125
L.b	Responses .....	126

# 1 Executive Summary — Part I

This Complaint document, in two Parts, alleges many very serious acts of misconduct, both corporate and legal in nature, perpetrated by Fritz Knabe, Dan Feldman, and others (IBM HR and Legal, personnel and processes), against myself (Walter Tuvell), during the period May–August, 2011. Part I deals with events in which Fritz was the dominant overt malefactor (with some aid from Dan); Part II deals with events in which Dan, together with allies in IBM HR and Legal took over as the main tortfeasors. Taken together, the two Parts tell a united, coherent, continuous story, and the two-Part document as a whole is intended to be considered and prosecuted as a single unit (though with multiple individual actionable elements to it). It is published in two Parts primarily for convenience of narration — two logically separable sets of activities are involved, and this organization greatly enhances understandability.

The structure of each Part is as follows:

- *An Executive Summary*, containing a “*list of particulars*”, providing a *summary* of the complained-of items.
- The *body* of the document, expanding the list of particulars into a *detailed narrative*.
- *Appendices* containing *supporting evidence*.

In addition, a single overarching *Theory of the Case* is given in Section 1.2.

Throughout, all dates are in 2011 unless explicitly stated to the contrary.

## 1.1 Disclaimer

For many years, I have used “legalistic” language on appropriate occasions (such as this two-Part Complaint, and in certain emails) — namely, when that language provides the most suitable communications protocol (most well-defined, descriptive, unambiguous) for the topic under discussion. However, I am only a hobbieist, not a “real” lawyer — such language is to be taken merely as a “layperson’s manner-of-speaking” only. No kind of “threat” (e.g., of “legal action”, or anything else) is to be imputed.

I hereby duly affirm I am fully committed to all policies/processes/procedures/practices supported by “IBM Law”,<sup>1</sup> as represented in the official BCG/AYJ/C&A documents. I have so affirmed from the very beginning of my employment at IBM/Netezza, and especially throughout this “unpleasantness”.

This does bring up a point right away, for those unschooled in the law of defamation and “rules of procedure (such as at court)”. Is there any sense in which what I have written in

1. By “IBM Law” (or, with a little sacrilege, “the Bible according to IBM”) we mean “the official IBM employee handbook,” a three-legged stool comprising BCG + AYJ + C&A (see Related Documents). The versions of these documents we rely upon (and quote) were retrieved from the IBM internal website, “w3”, in June. I first heard the phrase “IBM Law” from a presentation by Arvind Krishna, on Monday, March 7, when he used it at the “IBM/Netezza Corporate [Take-Over] Kickoff Meeting” (at Courtyard by Marriott in Marlboro), when he gave his advice about how to negotiate the IBM “system” using these words (essentially verbatim): “Don’t break financial rules, and don’t break the law — I mean the real law, not IBM Law — I don’t care about IBM Law.”

this Complaint (or in the emails cited/quoted in it, e.g.) can be called “defamatory”? No. Even beyond the absolute/unqualified privilege of pleadings, defamation requires “statement-of-fact, known false when uttered”. I hereby explicitly state: everything I have written here is either (i) statement-of-(believed-)fact based upon information/evidence, or (ii) statement-of-opinion based upon (believed-)information/evidence — in every case, believed (in good faith) true when uttered, as supported by reasoning stated herein.

## 1.2 Theory Of The Case: Blackballing

In a case as complicated as this one, involving so many people acting in close coordination, there is always an overarching theory motivating the culprits’ actions — one that explains everything in a comprehensible, unified way. It is difficult for the victim to discover that core theory, because it is a carefully kept secret amongst the co-conspirators. It has taken me many hard weeks to unearth and articulate the theory underlying this case, but I now feel confident I can formulate it in a coherent and correct way. The reason for my confidence is that *no other theory comes close to providing a comprehensive explanation*.<sup>2</sup>

Therefore, upon information, belief and evidence, I hereby allege the following “theory of this case,” as presented in the following two paragraphs, and detailed in the two Parts of this Complaint (these two paragraphs corresponding to those two Parts), together with the bold-face summarization sentence at the end of this Section:

S

IBM wanted to get rid of me because I am “too old”.<sup>3</sup> Age was not a barrier at the “old Netezza”, but it is at the “new Netezza”, i.e., at IBM.<sup>4</sup> Somebody in IBM HR learned I was the oldest employee<sup>5</sup> at Netezza, and raised it as an issue. (Almost-)legally, IBM could have terminated me “without cause”. But IBM’s lawyers thought that was “too risky” (i.e., IBM “couldn’t get away with it”), because I am a first-rate/model employee (as the record clearly shows), and ethical companies simply do not terminate such employees “without cause”. An “at-will no-cause” dismissal would have raised eyebrows — my eyebrows in particular. I am very experienced and intelligent, so IBM feared I would discover the real reason (ageism). I am also “too-trusting” (as are all purely technical PhD folks — and none are more purely technical than pure mathematicians), yet strong and fearless (due to experience gained from suffering previous untoward/illegal job actions).<sup>6</sup> That combination made me “dangerous”: IBM was worried I might file a lawsuit for unlawful dismissal. The publicity attending a messy high-profile lawsuit (especially one like this, with its (i) scandalous anti-societal message contrary to public policy, (ii) “big-guy vs. little-guy” overtones, and (iii) class action potential) could have spectacular long-term adverse effects for IBM: (i) erosion of public and partner confidence; (ii) difficulty recruiting competitive talent; (iii) plunge in prestige and value of the company (especially in this, its vaunted 100<sup>th</sup> year); etc. Therefore IBM decided I needed to be “blackballed”.<sup>7</sup> A pretended excuse for dismissal needed to be manufactured, one which

2· Occam’s Razor (*lex parsimoniae*): “Don’t multiply entities beyond necessity.” I.e., the simplest theory that suffices to explain events is probably the best/correct one (at least heuristically).

3· No other proffered reason has the “ring of truth” about it.

4· I was hired on November 3, 2010, exactly one week before IBM acquired Netezza.

5· That may not be strictly true, but I am almost certainly the oldest technical (non-managerial) employee.

6· “What does not destroy me strengthens me.” — Friedrich Nietzsche (*Twilight of the Idols*, 1888, Maxims and Arrows, 8: *Was mich nicht unbringt, macht mich stärker*).

7· (i) “Blackballing” is an informal/slang term for an illicit/unscrupulous form of harassment by management, whereby employees are “set up to fail”. E.g., onerous tasks are assigned which are impossible to fulfill (for various reasons, e.g., being overly difficult, insufficient time given, too vague,

could be plausibly defended in court (and IBM has been doing this a long time, with a lot of HR and lawyers skilled in this dark art).<sup>8</sup> My two co-managers, Fritz Knabe and Dan Feldman, were recruited to “do the deed”: Fritz would conduct an escalating campaign of not-so-subtly-false claims against me, hoping that would be enough to cause me to resign. But it wasn’t, so he ultimately had to explicitly pretend I hadn’t performed my job (though of course I had), and publicly broadcast that pretension by yelling at me in front of my peers. Dan thereupon dutifully demoted me — on the “basis” that I “couldn’t get along with Fritz” — in a very public way: adverse job-reassignment, from the highest-rung-on-the-ladder to the lowest. That was intended to have the emotional effect of embarrassing/depressing/scaring/coercing me into resigning, because I would “obviously” fear for my psychological safety, and I would certainly know my career at IBM was dead because nobody would want to work with me (I’d be “radioactive” due to the very public nature of the outburst and demotion). In short, the tactic was to intentionally instill sufficient emotional distress in me to cause me to resign: the tort of IIED (Intentional Infliction of Emotional Distress).

But I surprised everybody by not picking up on their “obvious” clues, and instead invoked “IBM Law”, IBM HR/C&A process. My “crime against IBM” was that I was “naive” enough to actually believe the “bullshit” “IBM written law” publicly promulgated to its employees through its “employee handbooks” (BCG, AYJ, C&A). I didn’t know the handbooks are in reality a sham, whose sole *raison d’être* is a publicity stunt — because that was a strictly kept “IBM unwritten rule”, so it could be used as a weapon to ferret out perceived “dissidents” (“not-team-players”). My invocation of the C&A process had the effect of irrevocably marking me as such a “dissident”. This made IBM boiling mad, so IBM doubled-down its efforts, now actively joining my management, HR and Legal together in the conspiracy. Through their combined efforts, though artless and undisguised to one as manifestly sane as I, their black-balling “game” became intense literally to the point of unbearability. Stupid “reasons” were conjured up, in an attempt to establish an absurd/twisted/inane/insane “paper trail” against me. Completely innocent utterances were pretended to be fraught with subversive meaning. Lies and inconsistencies were perpetrated, trying to “mess with my mind” (IIED again). “Rules” were (mis-)interpreted in the worst possible light to be applicable and/or enforceable against me solely, not others. Remarkably, I literally passed-out at a one-on-one meeting with Dan, due directly and precisely to the stress induced by his badgering. But even then I still didn’t “take the hint” and resign. Instead, I steadfastly followed the published IBM Law (instead of the “real/secret/surreptitious IBM Law”, of whose existence I was unaware), by having the “audacity” to take the “unheard-of” path of an Appeal. And when that turned out to be broken (incompetent, corrupt), I had to take the route of Confidentially Speaking and/or Corporate Open Door (“IBM courts of last resort”). The result is now this very two-Part Complaint.

etc.), and then the employee is terminated for “non-performance”. It is widely practiced (sometimes nameless to employees, but often known to HR personnel under the name “*managing-out*”), invariably kept secret from employees, and commonly used by unethical companies (but never by ethical companies). See especially the book *Corporate Confidential* by Cynthia Shapiro. For a brief popularization, see the *Reader’s Digest* article by Michelle Crouch (exp. p. 137).

(ii) The related term “*blacklisting*” also is well-known among HR/management. It is described this way by Michelle Crouch (online supplement): “17. ‘Companies do have black lists. It’s not written down anywhere but it’s a list of people they’d be happy to get rid of if the opportunity arises. If you feel invisible, if you’re getting bad assignments, if your boss is ignoring you, or if they move your office, you’re probably on it.’ — Cynthia Shapiro”.

8. “In some companies managing out can even include making an illegal firing appear legal. ... Most employees never see this one coming because in many companies it has become an art.” — Shapiro, pp. 3, 5.



In short, the theory of this case (“blackballing”, “managing-out”, “workplace-bullying”<sup>9</sup>) can be summarized as follows:<sup>10</sup>

**Premeditated conspiracy of defamation, deceit/fraud and IIED, with the aim of coercing me to resign — or alternatively to fabricate false “evidence against me” leading to “dismissal for cause” — all motivated by age discrimination.**

**Said conspiracy against me (and presumably many other employees) has been secretly sponsored/supported by IBM at corporate level for years, in the sense of systematic, widespread coordination involving management/HR/Legal.**

## 1.3 List Of Particulars

- On Wednesday, May 18, Fritz told Dan (according to what Dan told me) that I had disobeyed his (Fritz’s) orders by failing to produce certain Excel graphics for him. Subsequently that same day, Dan relayed that message to me. But there had never been any such order, explicit or implicit, direct or implied. Therefore either Fritz or Dan (or both, in some strange combination) was lying. If Fritz was lying to Dan, this constitutes defamation of me. If Dan was lying to me, this constitutes bullying/harassment/hostile-work-environment/IIED (Intentional Infliction of Emotional Distress) against me.<sup>11</sup>
- Motivated by the above “Excel graphics” event (and other events), I many times (at least six) asked Dan to call a three-way meeting (myself, Dan, Fritz), to clear the air (minimally, to come to a common understanding of the facts; maximally to smooth everything over so that normal work relations could resume). That request was consistently refused by Dan every time, either implicitly (by such meeting not being convened) or explicitly (by Dan’s actually telling me he would not convene such a meeting).<sup>12</sup> Moreover, Dan many times explicitly stated to me that his (Dan’s) “only [Dan’s word] concern was the success of Fritz’s project [code-named Wahoo]” — even after I pointed out to Dan that calling Wahoo’s success his “only” concern excluded any kind of justice for me. Both of these (refusal of three-way meetings, and “only concerned with success of Wahoo”) constitute a hostile work environment.
- On Wednesday, June 8, Fritz yelled/shouted at me very loudly, in public (open work area in Cambridge, with four co-workers present and observing). His reason for yelling, as he stated during his yelling, was that I had failed to do what he ordered/expected me to do (which was to measure some performance-reporting statistics during an overnight test run). But I had already informed him fully four times earlier that very day that the test run I had done was a “debug no-stats” run

9· See Wikipedia article on “workplace bullying” cited in Related Documents.

see also Add. III, p. 5

10· If this “conspiracy theory” seems far-fetched — it’s not. Conspiratorial blackballing is a rampant/common practice amongst unethical companies. For proofs and details, see Shapiro, *Corporate Confidential*. (Alas, I didn’t know this until August, 2011.)

11· It also constitutes defamation against Fritz. Legally, that is only actionable by Fritz, not by me. However, according to the BCG (p. 8), it is certainly proper for me to file a “third-party” complaint about it, and I hereby do so (if Dan did indeed lie).

12· As a clarification: I am here referring to “issues-related/procedural” three-way meetings, devoted to resolving “HR-type issues”. This is to be distinguished from “work-related/technical” three-way meetings, which are devoted to solving technical problems. (The latter type of three-way meeting did happen occasionally, but those meetings are generally not of interest for this document.)



(which is well-known at Netezza to be mutually exclusive with measuring performance-reporting statistics). Therefore Fritz was knowingly lying. This knowing falsity, in the presence of witnesses, constitutes defamation of me.

- On Friday, June 10, at a meeting in Dan's office, Dan told me that Fritz had told Dan that I was a "bully" and a "liar". That is wholly false, with no basis in fact whatsoever. Therefore either Fritz or Dan was lying. I have asked Dan multiple times what was the basis of Fritz's accusations, but he has consistently refused to tell me. This again constitutes defamation (by Fritz) and/or IIED (by Dan), because whoever was lying was doing so knowingly.
- On Friday, June 10, at the same meeting in Dan's office, Dan executed an adverse job action upon me, by removing me from my desired position, and indeed demoting me from the highest position in our Performance Architecture group to the lowest.<sup>13</sup> He stated no reason for doing so.<sup>14</sup> This (both the demotion, and the refusal to explain the demotion) constitutes a hostile work environment.
- The manner in which Dan demoted me was to "switch" me with Sujatha Mizar, a colleague of mine in our (Dan's) group, that is, to swap our job assignments. I am a man, far over forty years old, Caucasian. Sujatha is a woman, far under forty, non-Caucasian. Sujatha is also far less qualified than me (for example, I have a PhD while she doesn't, not to mention my decades of much more relevant experience).<sup>15</sup> This states a *prima facie* case of discrimination on the basis of age, and possibly also sex and race. This is especially so since Dan has refused to state to me any reason (much less explicit/coherent/truthful reasons) for the demotion, as noted above.<sup>16</sup>
- Due to the allegations of this Complaint, my reputation has been damaged (above and beyond the aforementioned defamation by Fritz), in the sense that I am now "blacklisted" ("considered radioactive") — not only by everybody who knows anything about my case (including the mere assertion of rights by invoking the

13. This constitutes a "significant tangible adverse employment action" in the sense of AYJ (p. 9), which explicitly includes "undesirable reassignment".

14. Dan did state that he thought Fritz and I "can't continue working together any longer", but that is merely a "conclusion", not a "reason (for the conclusion)". At this meeting Dan did also state that Fritz had called me a "bully and liar", though he didn't specify that was his reason for demoting me. It is, however, my belief that the latter ("bully and liar") was the "reason" for the demotion (leading to Dan to believe that was the reason for Fritz's actions against me, and to conclude that Fritz and I couldn't continue working together). The reason I wrote the word "reason" in quotation-marks in the previous sentence is that I allege age discrimination as the true overarching motivation behind this Complaint, but that "bully and liar" is the "made-up reason" for the job reassignment (i.e., the "[false]-documentation-reason", in the sense of "covering-up-tracks").

15. Nothing written in this Complaint (including cited emails) should be taken as criticism of Sujatha in any sense; I have no complaint against her at all. If any language seems overly harsh of Sujatha (or other "innocent bystanders"), that's not the intent — such wording serves only the requirements of complaint against IBM (Dan, Fritz, etc.), namely, the necessity that differences (of our stations in life/work, etc.) be duly recognized/emphasized, for reasons of "proper notice/argument".

16. That is: "If everything Dan was doing was on the up-and-up, why didn't he just tell me his reasons?" But even if Dan had stated to me explicit/coherent/truthful reasons for the demotion, the very act of the demotion would itself still constitute a *prima facie* case for discrimination.

see Add. IV, p. 8, middle bullet

IBM C&A process), but by many others as well, by “loss or reputation”.<sup>17</sup> This again constitutes hostile work environment.

## 2 Background — Part I

### 2.1 Introduction

After earning my BS degree at MIT, and PhD at the University of Chicago (both in Mathematics), I embarked upon a profession as an architect/designer/developer in the field of computer software, where I’ve had a long, diverse, and successful career of nearly 30 years, including a handful of patents. Before coming to Netezza, I worked at: AT&T Bell Labs (UNIX kernel); Open Software Foundation, OSF (distributed computing, security); Groove Networks, now owned by Microsoft (security); Juniper Networks (RADIUS authentication server); EMC, beginning in the VMware division before moving to the cloud-computing division (focus on performance).

I joined Netezza on Nov. 3, 2010, as a member of Dan Feldman’s Performance Architecture Group,<sup>18</sup> in Marlboro, Massachusetts. However, from the very beginning I was assigned “dotted-line” to Fritz Knabe, “embedded” in his Wahoo Group, in Cambridge. Dan and Fritz had both interviewed me during my interview cycle (as had John Metzger, their joint manager, and others).

Even though I have a great deal of industry experience, I had never before worked on a “database management system” (such as the Netezza Performance Server (NPS) product), nor had I “seriously” specialized in performance either.<sup>19</sup> These facts were of course well-known to everyone, and I was “up for” the challenge of “stretching” myself.

After an initial adjustment period, I settled into a routine of working four days per week on-site in Cambridge, and one day (Wednesday) in Marlboro. In this manner, both Dan and Fritz constituted my “equal co-managers”, though Fritz was somewhat “more equal than Dan” for all practical purposes (day-to-day operational management).

As a matter of working conditions, it should be noted that the Cambridge office was very small, housing only eight employees: Fritz, myself, Dan Dietterich, Rich Title, Devesh Agrawal, Steve Lubars, Jeff Keller, Huamin Chen (and also James Percent, but only for a brief period of 2-3 months). There, we all worked in a “public open-bullpen” environment, that is, a single big room with open (unwalled) cubicles, so that no privacy was available (there was also a meeting room with a door, and an unoccupied area of walled cubicles at the back of the office).

---

17. Note the legal principle applies in this case, that the defamation is “of a kind that tends to injure me in respect of my occupation” — and therefore no “actual injury” (“special harm”) need be shown. This acknowledges the ephemeral/unmeasurable nature of (loss-of-)reputation. Shakespeare (*Othello*, Act 3, Scene 3): “But he that filches from me my good name / Robs me of that which not enriches him / And makes me poor indeed.”

18. As opposed to, say, a “Performance Quality-Assurance Group”. It is my understanding that the idea of Netezza’s having an explicitly architecturally oriented performance group was Dan’s idea.

19. All software engineers are concerned about performance to some extent of course, but I had been labelled a “performance engineer” in only one previous job, though it was more “quality assurance” oriented, as opposed to “architecture/development” oriented, so that experience wasn’t as relevant as its title might seem at first blush.

## 2.2 Relations With Dan And Fritz

Before the difficulties described in this Complaint, I was aware of no conflicts of concern with either Dan or Fritz.

My relationship with Dan, in particular, I thought was an excellent one. I effusively told my wife many times that Dan was one of the best two managers I'd ever worked for. We connected with one another at many levels (personal, professional, intellectual). Like many new acquaintances, we did early-on have a few communication "impedance mismatches", but they were quite minor, and once identified/resolved were quickly patched over and caused no problems (certainly not in an actionability sense).

Very importantly for our relationship, Dan early recognized that I was "very shy" — pathologically so. That was a conscious behavior on my part. The reason was that I'd previously received some very poor workplace treatment previously. I was reticent about informing Dan of the reason for my shyness, but he was insistent, over many months (sympathetically, I thought, out of genuine caring), about "drawing me out of my shell". In particular, Dan was ultimately aware that: (i) I'd previously been subjected to hostile work environment, defamation, and "blackballing"; that (ii) I'd successfully taken legal action against the perpetrator of the defamation (against the company, not the individual; and it was resolved at the level of arbitration, not court); and that (iii) I really, really did not want to be involved in any "situations" at Netezza.

My relationship with Fritz was more complex. It was, for example, quite puzzling to me that even though Fritz was my manager 80% of the time (four days per week), he never really interacted with me directly in a normal manager/employee manner. His management style was dominated by a "scrum" methodology, so there were certain regular Cambridge team meetings: (i) daily fifteen-minute "stand-up" meetings of an update nature; and (ii) planning meetings for every "sprint/iteration" every two-to-three weeks. But Fritz and I never had "one-on-one" meetings in the normal manager/employee manner — even though he regularly held such meetings with all other Cambridge employees.

Consequently, I early had a sense of being a "second-class citizen (in Cambridge)". But I chalked that up to my own personal feelings, not to a failing on Fritz's part. Specifically, I deduced that I myself may have misconstrued my role, and that I really was in fact supposed to be a second-class citizen (in Cambridge). For, I was after all a member of Dan's performance group, "embedded" in the Wahoo group for the purpose of helping with the performance of the Wahoo product (as opposed to a more "architectural" role consistent with Dan's vision of a "Performance Architecture Group"). Coupling that with my relative inexperience in some relevant technology areas (as noted above), I thought that perhaps I just didn't understand what was expected of me (yet).

## 2.3 Wahoo Instability

Compounding my ambivalence about my role in Cambridge was the fact that the Wahoo machine (both hardware and software) was very much at prototype stage, and incapable of being performance-investigated most of the time. Only one Wahoo-class machine existed (in the "networking closet" in Cambridge), and the contention for it was great ("over-subscribed"). The Wahoo "developers" (that is, everyone in the Cambridge office but me) had priority getting access to the Wahoo machine (because the machine had to work functionally before its

performance characteristics could be seriously considered). And at the few times I did have access to the machine, it mostly didn't run at all (largely for software reasons, though sometimes for hardware reasons). All-in-all, when the machine did run, its performance was abysmally poor, unworthy of "serious" performance study (other than "mere measurement", that is).

The instability of Wahoo became a constant refrain (amongst the whole Cambridge team, not just me). My uneasiness with the situation (that I wasn't "pulling my weight") manifested itself in various discussions with Dan (recalling that Fritz unaccountably preferred not talking directly with me, as explained above).<sup>20</sup>

For example, the following is an excerpt from my weekly report to Dan for the week of January 30 – February 6 (Appendix A.l):

- Notably, didn't do much for Wahoo Perf. That's because Wahoo's problems are now at dev/impl/debug level, where I'm not currently qualified to help out. Without placing fault, this is generally a point of concern: am I being used, and/or contributing, the best way I can? Ongoing dialog with Dan about this.

Again, from my weekly report for the following week (February 6–13, Appendix A.m):

- Still unresolved is how exactly I'll be proceeding w.r.t. Wahoo. Late-breaking news at end-of-week is that Dan Dietterich has finally made the break-through on understanding where the cycles are disappearing to in Wahoo vs. Violin, so hopefully I'll be able to enter that arena again soon. Fritz also said he has some ideas in mind. Will find on [*sic: should be "out"*] early next week, when next "iteration" planning happens.

And again, from the weekly report for the next week, February 13–20 (Appendix A.n):

- Iteration 13 planning in Camb (Tue). Now that the weird Violin/vRAID perf issue has been resolved, it's time to do some more regular perf testing. It has been decided to starting [*sic: should be "start"*] with lower-level storage-layer testing, using DanD's PMtest (Partition Mgr), as opposed to the abortive PerfBar testing that was tried earlier (will get back to that later, when Wahoo is a little more mature). By end of week, got some PMtesting done. It's painful because Wahoo is still unstable, so the cycle is: (i) a few min of data collection; (ii) 20 min rebooting Wahoo and creating a test DB; (iii) lather, rinse, repeat, except that this can't simply all be put in a script because of unpredictable Wahoo hangs.

And yet again, from the weekly report for the week of March 13–20 (Appendix A.r):

- Tried to do another round of PerfBar/TPC-DS on WahooProto. Couldn't do it earlier, because either: (i) it was futile because WahooProto was known to be uselessly slow (before the "randomized scanlist hack"); (ii) it was over-subscribed due to merging; (iii) it would crash in the process of trying to build the TPC-DS database. Once those problems were fixed, I did a TPC-DS run, but it hung after ~20 queries. But those queries revealed another perf bug (still ~20 times slower than Skimmer).

My weekly reports are replete with many further comments about Wahoo instability/unusability, but the above few examples suffice to establish the point here. For more details (includ-

20. Fritz's reluctance in communicating with me was "unaccountable" at the time, though he finally expressed a "reason" for it at a meeting I had with Fritz on May 23 (Section 4.1).

ing the context in which the above examples occurred), see all my weekly reports (Appendix A). Those weekly reports provide a contemporaneous account of the events described in this Complaint (and Dan told me my weekly reports were unusually clear, interesting and enlightening).

## 2.4 Other Work

So instead of just “sitting back, waiting for Wahoo to be built”, I proactively sought to undertake some productive projects to “fill in my time”. Some of the resulting projects were Wahoo-related, and some were non-Wahoo-related; some were prompted by Dan or others (such as Garth Dickie and Michael Sporer); some were mostly invented by me. All were vetted by Dan, and Fritz was aware of all this too (I updated him daily at the Cambridge stand-up meetings, for example).

This resulted in two very significant accomplishments: (i) “PMtest” (noted above), a study of solid-state device I/O data rates (as implemented by Wahoo); and (ii) “PerfScore” (“Performance Scoring”), a new overall regime for comparing the performance of one NPS product/release against another:<sup>21</sup>

- The PMtest achievement was publicly held up on a pedestal by Dan (at a company-wide “Lunch-’N-Learn” presentation) as being the “Gold Standard” for performance-reporting at Netezza. Indeed, that had been one of the goals of the PMtest project, as established by Dan himself (he wanted me to write “a good example performance-report for other group members could emulate”).

change “could”  
to “to”

- The PerfScore achievement was even more notable, in that it introduced a level of proper scientific rigor to Netezza performance theory/practice, both numerical and graphical, that had previously existed only as muddle-headed hand-waving nonsense<sup>22</sup> — even though Netezza’s “core competency” is widely held to be performance *per se* (to the point that “performance” is explicitly NPS’s “middle name”).<sup>23</sup>

These two achievements were particularly notable considering my previous lack of experience in the performance arena (noted above). But these successes were even further augmented by several additional non-trivial achievements, going by such taglines as: (iii) “native Linux on laptop” (I’m the only person to succeed in doing that, despite attempts by several others); (iv) “build/development on Fedora (as opposed to on Ubuntu)” (ditto); (v) “CSV beautification tool”; (vi) “mugshot tool”; (vii) “probability distribution tool” (for a hashing study); (viii) “Netezza/IBM document template”;<sup>24</sup> (ix) “Project Glass” (with consultant Clark French; involved only short-term due to other commitments, Clark’s surgery, and broadened scope of

21. Note, very importantly, that the PerfScore project was actually prompted by Wahoo itself. Namely, my very first task for Wahoo was to “compare Wahoo vs. Skimmer”. But the existing method at Netezza for comparing any two NPS systems was woefully wrong, so I was prompted to research the field and establish the new PerfScore regime to implement a proper notion of comparison.

22. These derogatory words describe only the previous technical solution, and are not to be taken as an *ad hominem* slur against the people who instituted it (both because those people simply may not have had the necessary background, but also because I don’t even know who those people were).

23. Furthermore, the PerfScore achievement exceeded even what Dan himself had been able to accomplish when he himself had held a position in a previous company similar to the position I now hold. When faced with the same problem, the solution he came up with was essentially identical to the solution previously used at Netezza, until now superseded by my PerfScore solution.

24. This very document is written using said template.

see New Complaint, Add.  
III, bullet item on p. 12



his consultancy); (x) “spreadsheet model” (another short-term involvement, the lead developer being Dan himself); (xi) “NPS clock synchronization”; and others. Even following the difficulties described in this Complaint, my productivity has continued at a similarly high level, despite the projects being initially unfamiliar to me: (xii) “nzVtCapture.sh, capturing complete virtual tables” (utterly solving a problem that others had only fumbled at over a period of years); and (xiii) “Blktrace/Blkparse/Btt”, research project, ongoing at the present time, the ultimate outcome of which cannot yet be guessed.<sup>25</sup>

All-in-all, there can be no question whatsoever about the high level of competency, excellence, and mastery with which I have sought out and performed valuable work at Netezza, despite the balkiness of Wahoo.

### 3 Fritz: Database Transport (March 15-17)

Prior to the major difficulties described in this Complaint, I was aware of only a single incident that could even potentially be called anywhere near a “conflict” with Fritz. This seemed relatively minor at the time, for it was only I who was upset at the time (not Fritz), but in retrospect it should now be treated as a harbinger of the bigger problems to follow.

The incident occurred in the mid-March timeframe. The instigation for the incident was that Wahoo was misbehaving (as usual, see above), in the sense that on Tuesday, March 15, it was unable to create a test-database (“100GB TPC-DS”) that was needed for testing. The developers (Devesh, Rich, Jeff) knew what the source of the bug was, and how to fix it, and were in the process of fixing it, so the “right thing to do” was to wait until the bug was fixed (the next day), and then use the then-working Wahoo to create the required test-database. And in the event, that’s exactly what ended up happening (see the following paragraphs, and the email chain on this topic, Appendix B).

But although Fritz knew about the impending fix, instead of simply waiting for it to happen (and focusing on other, productive work in the interim), he demanded that I “spend” (really, “waste”) my time the next day, during my regular Wednesday (March 16) trip to Marlboro: (i) creating the necessary test-database on a different machine (in the Marlboro “co-location facility”, a few miles away); (ii) downloading it to my laptop; (iii) transporting it to Cambridge the next day (Thursday, March 17); and (iv) uploading it onto the Wahoo machine. This made no sense at all, for at least 3 reasons: (i) it was uncertain whether Fritz’s plan would work (because of the general instability of Wahoo, and this kind of uploading had never been attempted on Wahoo); (ii) the known bug would be fixed imminently anyway (so Fritz’s plan wouldn’t save anybody any time for the Wahoo project); (iii) I had much other, more important work to do (so Fritz’s plan would waste time I could usefully apply to other projects, especially the “Glass” project for Dan).

Fritz knew all these things, because I enumerated these objections to him; there’s an allusion to this in my email to Fritz, Appendix B 03/16/2011 10:45 AM, which Fritz never responded to. Nevertheless he was adamant. Therefore, I proceeded with Fritz’s plan (beginning on Tuesday evening by creating the test-database on the Marlboro machine). Here’s how I described the incident in my weekly report for the week of March 13-20 (Appendix A.r):

<sup>25</sup> There was yet another — very important — project that I was working on, ostensibly for Wahoo but with broader applicability for all NPS products, which was forcibly withdrawn from me and given to Sujatha. This was: (xiv) “PerfReport” (a.k.a. “WaltBar”).



- In the throes of problems (i)–(iii) above, Fritz sent me on a wild goose chase to ASCII-dump a TPC-DS DB from another machine, and hand-load it into WahooProto (as opposed to regular backup/restore, because WahooProto doesn't support restore yet). I was upset about this, because I was already over-worked trying to finish up PMtest and Glass. As a result, I had to relinquish my participation in Glass. Sad day.

The “sad day” referred to here was the relinquishment of my participation in Glass. It was an interesting project and I wanted to continue with it, but Dan decided to remove me from Glass, for reasons stated in the next Section 3.1.

### 3.1 Aftermath

In discussion soon afterwards with Dan about this incident, he informed me that Fritz “seemed to be feeling a bit uneasy” (paraphrase) about the level of my participation in projects not directly related to performance-testing of Wahoo. I stated to Dan that there was nothing I could be doing with Wahoo that I hadn't already done (because of Wahoo's instability/unusability, citing reasons as explained above), so I was filling-in my time doing projects that had valuable tangible benefits, and at no time was I slighting Wahoo work in favor of non-Wahoo work. I emphasized to Dan that never at any time during my participation in the Wahoo project have I ever slighted the concerns of Wahoo in favor of non-Wahoo projects. Period. Dan agreed with me of course (after all, he already knew this all along). In any case, Dan sought to assuage Fritz's concerns (among other reasons, as noted above, as Dan told me) by unilaterally removing me from the Glass project, thereby demonstrating good faith to Fritz about devoting my time more fully to Wahoo (note: Dan had also assigned me to work on Glass unilaterally, without asking me what I thought about the matter, but I view both assignment/unassignment as within his purview as my manager).

Notably (as part of a continuing bad pattern), it is to be observed that this discussion took place only between Dan and myself, with no direct involvement from Fritz (as noted above, Fritz had an unexplained aversion about interacting directly with me in a normal manager/employee manner). Any/all discussions about the matter between Dan and Fritz had been done between the two of them only — without my presence or knowledge.

Nevertheless, at my request, a three-way meeting (me, Dan, Fritz) was held in Cambridge (I can't find a contemporaneous record of it, but my recollection is that it happened on Friday, March 18). There, I repeated the arguments I'd given to Dan (see first paragraph above in this Section), with Dan agreeing with me, and Fritz then also recognized the truth of these assertions, and he agreed too. So, I thought that would be the end of any potential friction with Fritz.<sup>26</sup>

## 4 Fritz: Excel Graphics (May 18)

Everything was fine with me and Fritz (and Dan), I thought. The first inkling I had that serious trouble might be brewing with Fritz (as opposed to the seemingly minor incident of data-

26· It was the success of this three-way meeting that gave me some confidence that subsequent three-way meetings could be instrumental in resolving issues between myself and Fritz. Nevertheless, when I requested such three-way meetings to Dan later (see below), they were denied, for reasons unknown to me.

base transport, above) was the major blow-up described in this Section. It occurred at a one-on-one meeting I had with Dan on Wednesday, May 18. The story requires some set-up.

Towards the end of March, the Wahoo machine finally did start becoming useful for performance-testing purposes, and the Cambridge team began designing appropriate performance test plans. It was my task to implement the plan, piecemeal, step-by-step, collaboratively with the team (as is normal in the software development business, and was expected on the Wahoo project).<sup>27</sup> This effort proceeded in nominal fashion through the month of April, and into May. It was, for example, incorporated into Fritz's scrum planning, and it was discussed daily at Fritz's stand-up meetings. At no time was there any dissension voiced by anyone. The plan was to write a set of shell-scripts that would run performance tests on Wahoo, and generate an informative ASCII report document, ultimately including an "ASCII-art" graph of certain important performance characteristics. I dubbed the whole scheme "PerfReport", though some team-members later started calling it informally "WaltBar" (in homage to my name combined with the name of a pre-existing performance testing scheme called "PerfBar" ["bar" for "clearing the 'bar' {or 'hurdle'}"]).

9, not 16, see Add. V, pp. 21-22

On Monday, May 16, following the usual Wahoo stand-up meeting (where I reported my status with the WaltBar work), Fritz was reviewing WaltBar at his desk, and called me over to discuss it.<sup>28</sup> He suggested that he'd like to see, as part of the WaltBar report, an indication of resource capacity-usage, that is, where demand on Wahoo resources (mainly CPU/FPGA capacity, and Violin storage device bandwidth) started nearing their limits ("becoming over-subscribed"), and he suggested a figure of 80-90% for this purpose (stating that it was my choice). I immediately told him this was a great idea, and that an improvement would be to construct an ASCII-art graph (mentioned above), in such a way that the reader of a WaltBar report could visually detect an over-subscription scenario, and could decide for themselves their own value for over-subscription (as opposed to our deciding in advance the 80-90% figure Fritz suggested). He agreed, and I immediately set to work implementing the suggestion.

The next day, Tuesday, May 17, Fritz called me over to his desk again, late in the afternoon (about 4:00 PM) and we had some further discussion about WaltBar. I had spent the day implementing the ASCII-art graphic discussed the previous days, and I felt very good about it, because it was turning out to be a very informative reporting technique. Fritz's thoughts during this discussion were wide-ranging, exploratory and vague, unlike the preceding day where he'd made a specific suggestion (and I implemented it). For example, he showed me, on his wide-screen Apple Macintosh monitor, how he could use Emacs to re-format some of the WaltBar reporting tables (from "plain ASCII" to "comma-separated values (CSV)"), because he said he was interested in doing some additional analytical investigation on the numbers, making some cryptic scribbles on a small (3"×3") Post-It note. I offered to reformat some of the WaltBar reporting tables to make his job easier (eliminating the Emacs re-formatting), but he said that wouldn't be necessary. Indeed, nothing Fritz said amounted to an action-item for me of any kind, and with that I left the office for the day.

27. This kind of activity was not something I could accomplish all by myself, for example, because of my inexperience with both "architecture-level" performance-specific testing, as well as "knowing what was already known via the collective wisdom of Netezza development in general, and the Wahoo development team in particular". Nor was that expected of me, because this sort of activity involves an iterative process of trial-and-experiment, or "method of successive approximations", by the team as a collective whole.

28. This was in itself unusual, because as noted above Fritz rarely spoke directly to me. But it still wasn't a normal "one-on-one" because it was still in the public open-bullpen environment of the Cambridge office.

On Wednesday, May 18, I reported for work in Marlboro, per my usual weekly schedule. There, I attended Fritz's regular stand-up meeting at 11:00 AM (Fritz and some other members of the Cambridge team also had a habit of working at Marlboro on Wednesdays, for various reasons), followed by Fritz's regular Wahoo project status meeting at 11:30-12:00 (attended by the Cambridge team, as well as several others, including Dan, Michael Sporer, Garth Dickie, etc.). At the latter meeting, Fritz presented some Excel graphics, which depicted the same data as my ASCII-art graphics.<sup>29</sup> I recognized that Fritz had spent significant time since the previous day working on his graphics, however I and others recognized Fritz's graphics as representing nothing more than data collected by the WaltBar tool (several others even nodded in my direction as a congratulatory gesture, so I felt very good about that).

Later that afternoon, at 3:00 PM, I had my regular one-on-one meeting with Dan, in his office. I carried with me two pages of paper I'd just printed out, representing my latest work on the PerfScore project, feeling very good about it, and anticipating a nice "show-and-tell" for Dan. So I told Dan I had something to show him. But Dan also told me he had something to tell me, and he asked me which I wanted to do first. I had no idea what Dan had in mind (it was very irregular for him to surprise me with something unannounced like this), so I said "I know what mine's about, but I don't know what yours is about, so why don't you go first". And with that I laid my two pages of paper face-down on his desk (planning to pick them up again when it was my turn talk).

And then came the bombshell, totally out of the blue. Dan told me: "Fritz is ripping mad at you for disobeying his orders, by not producing the Excel graphics for him he asked you do do!"<sup>30</sup> I was utterly thunderstruck, and struck dumb, unable to reply for a minute, for I had absolutely no idea what was going on. Fritz and I had met for only a few minutes the previous day, late in the afternoon (certainly too late for Fritz to have asked me to any work for the next day [he always respected personal non-work time]), and Fritz had asked me to do nothing at all, much less produce Excel graphics. I am certain about this, since I surely would have remembered any such request without fail — *because I don't use Excel* (I don't even use Windows or Mac: I use open-source alternatives, Linux and OpenOffice.org Calc, and Dan and Fritz knew this).<sup>31</sup>

29· Without meaning to "dis" Fritz's work, I will remark that the kind of Excel graphics he produced are generally considered harder to interpret than the ASCII graphs I produced (because his graphs lumped four different resource-usage timelines onto a single graph, whereas WaltBar kept them distinct). I'd considered both kinds of graphs the previous day, when I adopted and expanded Fritz's original suggestion along these lines, and adopted the ASCII graphs for this reason of easier interpretability.

30· This is a paraphrase, but is an accurate portrayal of the impact Dan's words had on me at the time. The actual words he spoke may have been somewhat different, but I was so shocked I cannot now recall his precise wording. Later (at a one-on-one meeting sometime during the following 3 weeks, whose exact date I cannot precisely recall), Dan tried to water down what he had told me at this time, by pretending he had told me merely that Fritz was "frustrated" that I had not "picked up on Fritz's suggestion" that I create the Excel graphics. But that is an invention, not an accurate reflection of what Dan really told me. Indeed, this watering-down of Dan's did not occur the first time I brought up this incident to Dan, but only the second or third time I did so.

31· Excel does not run on Linux natively. The only way to run Excel on Linux is via a virtual machine (such as VMware or VirtualBox), but even that involves obtaining licenses for both Windows and Excel, and installing/initiating them, then using them to generate the desired graphics. Alternatively, OpenOffice.org can emit documents in Excel-compatible format, however that transformation is not 100% accurate, so had I thought about that, I would have told Fritz about potential incompatibilities, noting that we'd have to exchange the generated graphics and he'd have to check to make sure they worked on Excel before attempting to present them. All this is speculative though, because Fritz didn't hint he wanted me to do anything.

change "the  
previous day"  
to "earlier"

When I regained my speech, I tried explaining all this to Dan, but he seemed “deaf to my story” (specifically, his demeanor seemed condescending to me, as if he believed Fritz’s story wholeheartedly, and was only humoring me by listening to my side of it). As I tried to make myself truly understood, I became more and more frustrated/upset, as it became more and more clear Dan wasn’t really interested in listening to what I had to say. The more I tried, the more Dan stonewalled. It got to the point where I said words to this effect: “What Fritz said was factually false, but I hesitate to call it ‘lying’, because to my mind that word implies ‘intent to mislead’, but I’m not sure I can make that claim — for example, maybe Fritz is just under a lot of stress and expects me to be a mind-reader.” But, reminding Dan of the earlier database transport incident (described above), I did tell him that I now thought Fritz was a “workplace bully”.

At some point, approximately 20 minutes into our one-on-one, I recognized I was getting nowhere, and becoming really rather angry (incensed, defamed, “mad as hell”) about my plight. It became to bed that I stood up, and took a step towards the door, as if to “run away” (I was sitting between Dan and the door, so this motion led me to turn my back on Dan). But after that first step, I regained my senses (realizing that such an abrupt exit would be “unprofessional”), and stopped in my tracks, and turned to face Dan, and asked, “May I leave the room?” He said, “Yes.” So I turned again towards the door. But as I was just inches from grasping the door handle, Dan said, “What about these?” I turned again, and saw that he was holding the two pages of paper I’d brought into his office, extending them towards me with both his hands. At that, I said, “Here, I’ll help you.” I then took the papers from his hand (gently, as he was already handing them to me), and I turned slightly and threw them (silently and emphatically, but not violently or threateningly) into the paper recycling bin in Dan’s office.<sup>32</sup> With that, I silently strode towards the door again, and this time I reached it, opened it, and stepped out. Dan stepped to the doorway too, and called after me, “Where are you going?”, and I said, “Home”. Then he said, “Are you coming back?”, and I said, “Not today.”

And therewith I left the office for the day and drove home. It was 3:30 PM.

## 4.1 Aftermath

The next day, Thursday, May 19, I was still very upset, too upset to work near Fritz in Cambridge, so instead I worked-at-home.<sup>33</sup> The email chain revolving around this work-at-home day is recorded in Appendix C. As that email chain shows: (i) I was still quite mad about how I’d been treated; (ii) I requested Dan’s help in resolving it; and (iii) when he asked for input from me on how to proceed I suggested a three-way conversation (me, Fritz, Dan) as a way forward (similar to the one we’d had after the database transport issue, see above).

On Friday, May 20, I also worked-at-home, still too upset to face Fritz. My weekly report for that week reflected my upset (Appendix A.aa).<sup>34</sup>

32· This was not premeditated. Rather, I just happened to notice the recycle bin just as I took the papers from Dan, and spontaneously used it (because I normally use recycle bins when discarding paper).

33· I also had two doctor visits scheduled for that day, and as well I had earlier asked for that day and next Thursday and Friday, May 19-20, off so I could do some housework on my deck. But it turned out the carpenter was sick on Thursday, so I worked-at-home for Netezza instead.

34· That weekend, May 21-22, is when the carpenter was well enough for him and I to work on my deck, so I didn’t actually have to take any time off of work for that personal activity.

I can't locate any other contemporaneous records for this period, beyond my weekly report (Appendix A.r).<sup>35</sup> So the following dating might be a little off, but my recollections about the events seem clear enough to me: "Thank you SameTime" found later, New Complaint, Add. III, pp. 6-7

On Monday, May 23, I reported to Cambridge for work, and Fritz was there too. During the morning I was still stewing about the Excel graphics incident, but was unsure about how to handle it. The problem for me was that Dan had explicitly denied my request for a three-way conversation (see above). And, as has already been mentioned, Fritz continually failed to initiate conversations with me. Eventually, I decided I needed to initiate a conversation with Fritz, so I sent him an instant message and email (even though he was right there in the office, a few feet from me), and requested a one-on-one meeting with him. He accepted, and we went into the conference room. I began our meeting by telling him that Dan had told me he (Fritz) was upset about my not creating the Excel graphics for him the previous week, but I didn't recall any such request, so I wanted to know what that was all about. Fritz was evasive, giving no satisfactory answer to my query, but did acknowledge that my account of the incident (above) was correct in all its essentials (in particular, his expression of a vague desire to analyze the data more closely, but no request to me to create Excel graphics for him). The meeting ended after a short time (approximately 15 minutes), with me stating to Fritz, "if there are any further communication issues between me and you (Fritz), I'd appreciate hearing about it before Dan did". He agreed that was a reasonable request, and he would do so. After I explained some of the technical barriers that I'd been meeting all along that prevented me from making more/faster progress on Wahoo performance work (rather similar to those documented in Appendix E), Fritz acknowledged that he perhaps hadn't had a "proper perception/appreciation of the work involved" (mentioned in my weekly report, Appendix A.r). So he suggested that he would like to have more direct control over my time (as opposed to my remaining a "direct" report of Dan's, and only "dotted-line" to Fritz). I agree of course (indeed, I didn't even understand until that instant that Fritz somehow didn't feel "empowered" about owning my time, as he claimed he didn't — I'd understood all along that Fritz was really my "operational day-to-day boss", and Dan was just my "organizational/org-chart boss").

"vague  
desire"  
were  
Fritz's  
own  
words

"agreed"

On Tuesday, May 24, Dan visited Cambridge, and he and I had a one-on-one meeting there. I told him that I'd initiated a meeting with Fritz to discuss the Excel graphics incident (see above). But Dan continued to deny me the three-way meeting I'd requested. Instead, Dan suggested that the best way to move forward was for me to keep a day-by-day written journal (emails) of work and communications between myself and Fritz, so that Dan could keep closer tabs on our situation. I took some offense at this, stating that it seemed like Dan was mistrusting me, and placing me on some kind of "performance review". Dan denied that however, instead stating that his motive was to monitor Fritz more closely, so that he (Fritz) couldn't make future unreasonable claims about what he asked me to do. I didn't really believe this, but I had no alternative but to just "play along". So I told Dan I'd comply.

"move"

Towards the end of the one-on-one meeting, Fritz joined us, for the specific purpose of requesting Dan that he (Fritz) be allowed to have more direct control over my time (as stated above, during my meeting with Fritz the previous day, Monday). (I knew Fritz was going to join us, because I'd mentioned to him that Tuesday morning I was going to have a one-on-one with Dan, and Fritz asked if he could join us towards the end of the meeting to broach this question, and I agreed.) Dan agreed to the plan, as Fritz and I already had. Even after this formal agreement, however, it transpired that in the event Fritz did not change his behavior

35. Of course, access to IBM's document-retention databases would help here. I'd welcome that.



towards me going forward one iota (in particular he did not communicate with me more often, nor did he hold any one-on-one meetings with me).

It was also sometime later the week of May 22-29 (I think it was Thursday or Friday, May 26 or 27) that Dan proposed a solution to a certain technical problem we were having, by making a suggestion to me and Fritz that I look into using an open-source program, “nmon” (“Nigel’s monitor”), for the purpose of “host and blade monitoring”. I agreed to look at it, and immediately did look at it, and quickly discovered it was “almost exactly what the doctor ordered”. Specifically, nmon’s only shortfall was that it only supported data collection at the rate of once per second, but we needed it at the rate of ten times per second. So I set about modifying nmon to suit our needs (renaming the modified code “nzmon” in the process, to avoid confusion). See Appendix D, as well as my weekly report for May 29 – June 5 (Appendix A.cc).<sup>36</sup>

## 5 Fritz: Wahoo Not The Bottleneck (June 6)

As noted above, beginning on May 23 I started complying with Dan’s suggestion of keeping him and Fritz very closely updated day-to-day about the work I was doing for Fritz; this consisted of trying to performance-test Wahoo, but being blocked by Wahoo’s “non-compliance” (i.e., it was too broken to be performance-tested). This non-compliance of Wahoo was particularly aggravating for me (more so than for anyone else), because it was I myself who bore the brunt of Wahoo’s unworkability, in two ways: not only did it mean that (i) I couldn’t do the needed performance-investigation of Wahoo; but also that (ii) I was the very person who was in charge of trying to get Wahoo working, but it kept failing, thereby wasting very many of my hours (which could/should have been more usefully put to doing other things, especially developing WaltBar/nzmon).<sup>37</sup> In any case, my constant updating Dan and Fritz about Wahoo’s recalcitrance is evidenced by this comment in my weekly report for the following week (May 29 – June 5; Appendix A.cc), wherein I referred to the long-continuing breakage of Wahoo this way:

- Updated Fritz and DanF far too often, “as if” this weren’t a typical week in Wahooland, but it was.

It is rather tedious to submit more evidence than this comment for the ongoing balkiness of Wahoo to being performance-studied, for Wahoo’s performance problems are very well-known and disputed by no one. Nevertheless, for the sake of completeness, such evidence is indeed presented for this period. Appendix E.

Nevertheless, despite all this tedious and well-known evidence (all of which Fritz well knew about), on Monday, June 6, Fritz made a very remarkable verbal comment to me. That morning, Fritz and I arrived at the Cambridge office early, before anyone else. After exchanging pleasantries, we were sitting at our respective seats, working away, when he abruptly blurted

36· This nmon/nzmon work is mentioned here, not really because it represents a significant effort (though it does), but primarily because it plays an incidental role in the story, below.

37· In the vernacular of the software industry, my attempts at getting Wahoo to run made me the *de facto* “Quality Assurance (QA, i.e., ‘testing’)” team in charge of Wahoo, and the continued failures of Wahoo to run meant that the “Development team” (the other members of the Wahoo team in Cambridge) were “screwing QA by dumping bad software on me (‘throwing it over the wall’), thereby wasting my time trying to test an untestable system”.



out words substantially to the following effect (seemingly after reading all his email and assessing the current state of the Wahoo project):

- “Even though Wahoo still isn’t running (so it can’t be performance-tested), it also seems like you’re not getting your own work on nmon done either, so it looks like Wahoo is not the bottleneck [*that is, you are the one who is the bottleneck*].”

That was a very stupid, hurtful, and false thing for Fritz to say about my productivity. For, he knew full well (due to my “far too often” updating Fritz and Dan, noted above) that *the entire fault for schedule slippage was the development team (not me), because they kept dumping bad Wahoo software on me, thereby wasting my “QA”-time, and preventing me from doing my work on either performance-testing or on nmon/nzmon.*

Because of this idiotic, intentionally hurtful and harassing comment by Fritz, I was too steamed to be alone with him in the office. So I walked out the Cambridge office door, on the way out mumbling something along the lines of: “You may not think it takes a long time to get Wahoo working, but it does”. I walked around Kendall Square (where the Cambridge office is located) aimlessly for awhile, returning to the office only after 45 minutes, when I was sure other Cambridge team members would be in the office (as they were when I returned).

Because this comment of Fritz’s was verbal and nobody else was present, the first contemporaneous written record I have of it is embedded in the email I sent two days later on June 8 (Appendix I.a),<sup>38</sup> where I called his comment “an absurd statement”.

## 6 Fritz: Yelling In Public (June 8)

The final straw in the relationship between Fritz and myself was the yelling incident of Wednesday, June 8.

Paradoxically, the day began on a note of high euphoria for the whole Wahoo project (“there was singing in the streets of Cambridge,” is how I put it in my weekly report for the week of June 5–12; Appendix A.dd). The Wahoo machine had, finally, stabilized to the point where it could be performance-tested. This was the point we’d been trying to reach for two-and-one-half weeks, so I seized the opportunity to run the required test (“PerfBar TPC-DS@100” — this had been our standard, touchstone test for many months, and was still so at that time, see Fritz’s note about “next steps”, the first email in Appendix F) the night of Tuesday, June 7. When I arrived for work in Cambridge early the next day (Wednesday, June 8), I “ran the numbers” (that is, compared Wahoo’s performance to Skimmer’s, using my PerfScore mechanism), and was very surprised to discover that Wahoo’s speed came in at better than “4×” the speed of Skimmer, even though the test had been done using the “debug build” of Wahoo (which is slower than an “optimized/turbo build”; using a debug build also implied that “no statistics” were gathered, because gathering statistics is unreliable when using a debug build). After so many long months of failure, here was a welcome success indeed!

for much more detail about this testing,  
see Add. V, Sec. 57.6, pp. 24-28

<sup>38</sup> The main thrust of the June email 8 was not this “Wahoo not the bottleneck” incident, but the incident of Fritz’s yelling at me, described in the next section.

I immediately wrote-up and printed the two-page report (reproduced here in Appendix G, consisting of a graphics page and a numerics page).<sup>39</sup> I communicated this happy news explicitly to Fritz (and others in the office) no fewer than four times that morning, before noon:<sup>40</sup>

- As soon as I had written-up and printed the report (a little before 9:00 AM), I stood at the front of the room and publicly announced the result verbally.
- At the same time, I also posted a hardcopy of the report publicly on the whiteboard in the Cambridge office. At the time of posting, Fritz came up to me at the whiteboard, and we discussed the report. He fully understood its meaning, and he in fact actually corrected a small misstatement I inadvertently made about it (regarding the meaning of the “red line” on the graphic).
- I emailed the report to everyone on the Wahoo team, and to Dan, and 9:06 AM.
- I reiterated the same result yet again at Fritz’s regular 11:15 AM stand-up meeting.

In each of these instances, I stated “better than 4× faster for PerfBar TPC-DS@100,” while always explicitly emphasizing “debug build, no stats” (for example, these things were explicitly stated in the report itself, hence also in the email containing the report I sent to Fritz and others). Everyone’s attention was intensely riveted to this wonderful news, especially Fritz’s. There can be no question that this information was perfectly well understood by everyone, especially Fritz (as evidenced by his various comments over the weeks/months about performance-testing).

Also during the stand-up meeting, Steve and I spoke about our progress on the “Fudge Factor Model (FFM)” that Steve was doing for Fritz. This FFM consisted of some Excel graphics (the same graphics Fritz had generated earlier, with new input data), together with some numerics involving “fudge factors” to account for differences between the prototype Wahoo machine and the ultimate completed machine (both graphics and numerics based on WaltBar/PerfReport test results that I was to supply to Steve). What Steve and I reported was that we were still working out some protocol/formatting issues between ourselves, but that we were nearing the end of that task, and we committed to presenting Fritz with “at least one” example (namely, for the EXCEPT test from the Atomics suite) by close-of-business that day. All this was according to Fritz’s plan as expressed in his “next steps” email (first email in Appendix F), and he indeed seemed pleased by it.<sup>41</sup>

After the stand-up meeting, Fritz disappeared from the office for awhile (perhaps to go to lunch and/or attend a meeting or run an errand, I don’t know, and in any case it was not unusual for any of us to disappear for awhile like that). He returned later in the afternoon, at approximately 2:45 PM, and went to his desk (probably to read email, as would be normal). Then, Fritz approached Steve’s desk-area (which was next to mine), and leaned against it (partially sitting on it), and started talking to Steve. This was also not really very notable, and I ignored it (as did everyone else in the office). Steve’s seat in the Cambridge office was next to mine, and Steve was seated in it (with Fritz to his right and me to his left), so I no-

39. Also included in Appendix G is an email chain, explaining that the original report had been in error. That, however, was irrelevant to the yelling incident, because that report of error occurred hours afterward.

40. In contemporary emails of that period, I spoke of “three times”, but that was erroneous, because I neglected to count the hardcopy posted on the whiteboard as the fourth time.

41. The results of the Fudge Factor Model were to be included in some material Fritz was to present at a customer meeting immediately following Netezza’s “Enzee Universe” conference, scheduled for June 20-22 in Boston (i.e., some 2 weeks in the future), so there was no real hurry to get the work done.

ticed “out of the corner of my eye” their conversation was happening, but I paid scant attention to it because I was absorbed in my own work.

But at approximately 3:00 PM, at some point during their conversation, I heard my name mentioned, which piqued my attention, so I nonchalantly glanced at them and started listening (this being an “open bullpen” work area, that kind of thing is quite common). It transpired that they were discussing the “Fudge Factor Model” (mentioned above). As I listened, I observed that Fritz was “quizzing Steve down” about when Steve would be supplying Fritz with the FFM graphics and numerics, and Steve was saying that I had not yet supplied him (Steve) with the required WaltBar/PerfReport results. So naturally I joined the discussion, supporting what Steve was saying, and reminding Fritz that Steve and I had already reported during the stand-up meeting that morning the status of the FFM work, and reiterated that we’d be delivering the Atomics EXCEPT example to Fritz by close-of-business.

At that point, Fritz turned away from Steve, and directed his attention squarely at me, and became very animated, and angrily attacked me by *loudly yelling/shouting* at me (in full public view of Steve and the others in the office, Rich, Jeff and Huamin)!<sup>42</sup> Fritz was yelling that he expected me to have already supplied Steve with the WaltBar/PerfReport results he needed to generate the FFM results. But that was of course impossible, and Fritz knew it, because I’d already reported *four times* that morning that the results I’d run overnight were: (i) PerfBar (*not* WaltBar/PerfReport, as Steve needed); (ii) debug (*not* turbo, as Steve needed); and (iii) no-stats (*not* stats, as Steve needed). Furthermore, Fritz could not have expected me to have done the necessary test run (WaltBar/PerfReport, turbo, stats) during the period when he was out of the office that day, because: (iv) Steve and I had only committed to supplying him with the single Atomics EXCEPT example by end-of-day (*not* the entire TPC-DS suite, as Fritz wanted); and (v) running the entire TPC-DS suite on WaltBar/PerfReport would have been impossible in the time allotted (i.e., the time Fritz was out of the office), because that took approximately 7½ hours (as opposed to 1½ hours for a similar PerfBar run), and Fritz knew it. At some point during his yelling, Fritz said something to the following effect: “I know you said you weren’t running stats, but I thought you meant you *weren’t* running ‘PerfBar/PerfScore’-stats, but that you *were* running ‘WaltBar/PerfReport’-stats”. When he said that, I really got confused, because it made no sense at all (and Fritz knew it well, for Fritz is a micro-manager concerning technical details, and he’d previously spoken to me directly, in detail, about the contents of both PerfBar/PerfScore and WaltBar/PerfReport) — but under the circumstances (of being attacked), I didn’t realize how nonsensical it really was (namely, it’s nonsensically impossible to expect to get “WaltBar/PerfReport stats from a PerfBar/PerfScore run”), so I wasn’t able to give him the “perfect come-back”.

I was of course in a state of utter shock to find myself under siege by Fritz’s yelling this way (I’ve never seen it done anywhere before, nor even heard about it being done to anyone, much less to me). As he continued (for a upwards of a minute [it seemed like more]), I eventually found myself unable to stand it anymore, and I involuntarily said the exact words, “Fritz, get off my back” — in a voice louder than my usual volume, because I had to make myself heard over his yelling, but no louder than Fritz’s normal, non-yelling, volume. This was my only “in-anger” contribution to the entire yelling incident, but it was certainly not “aggressive”, it was entirely “defensive”.

---

42· It is to be noted here that Fritz has an extremely loud voice even under normal circumstances, so his yelling was a noteworthy event indeed, and could not have been missed by anyone anywhere nearby except the deaf.

Once I had “stood up to him” in this way, Fritz said (still quite loudly, and obviously derisively), “Well, I can see that verbal communication with you just isn’t working, so I’m going to have to write everything down from now on.” Thereupon he stalked back to his desk and started typing at his workstation — presumably to work on his email concerning “upcoming performance tests”, which I received shortly thereafter (Appendix H).

At this point I was of course very shaken. I wanted to “run away” (i.e., leave the office and go home), but I concluded that would seem “unprofessional” (similar to what I’d thought about at the one-on-one meeting with Dan on May 18, concerning the Excel graphics incident, see above). So I turned to my laptop (thereby turning my back on everyone else in the office), and tried to get some work done, but I couldn’t concentrate. After several minutes of calming down, I left the office and went home in a normal manner. The time was approximately 3:45 PM.

## 6.1 Aftermath

I went home from the yelling incident in somewhat a state of shock, for obvious reasons. That evening, I composed myself to the point where I was able to write-up my account of the incident, and mail it to Fritz and Dan (Appendix I.a).

Later that night, I also discovered, much to my amazement and chagrin (because I am uncommonly careful and bug-free), that my report of Wahoo being 4× faster than Skimmer was in error. So I also immediately wrote that up too and emailed it that night, with a follow-up the next day (both included in Appendix G).

Dan did more than “ask”, he pushed; after (i) he pushed me to come up with a second reason; then after (ii) he pushed me to come up with a third

The next day, Thursday, June 9, was a quiet one. I worked in Cambridge, while Fritz went to Marlboro (per prearrangement, because Arvind Krishna was visiting Marlboro that day). Late in the day, Dan (in Marlboro) called me by phone, and we talked about the yelling incident. Dan asked me if I could think of why Fritz had acted the way he did. I told Dan I could think of only three possibilities, which I told to him: (i) Fritz was under a lot of stress (to demonstrate major performance improvements in Wahoo, because it wasn’t living up to his claims for the project); (ii) Fritz was “threatened by me”, though I had no idea why; (iii) Fritz was going insane (literally, but also this point was a “catch-all” for all other reasons, as I told Dan). Dan explicitly rejected (ii) and (iii) out of hand, leaving us with only (i) to work with. Dan told me to try thinking about ways I could “patch over my difficulties with Fritz, and make him happy” (paraphrase). I also requested to have a three-way meeting amongst the three of us, but Dan refused to grant that request.<sup>43</sup> At the end of the phone call, Dan asked me if I could travel to Marlboro the next day (Friday), and have a one-on-one meeting with him to continue our discussion, and I agreed to do so.

On Friday, June 10, I had my one-on-one meeting with Dan, first thing in the morning. There, Dan disclosed to me that Fritz had met with him the preceding day (Thursday), and that it was “now too late to try patching over the difficulties, Fritz”. Instead, Dan said he’d decided that Fritz and I could no longer work together. I told him I was thinking the same thing — but I didn’t get to say another word,<sup>44</sup> because Dan immediately “gave me the bum’s rush”, i.e., started speaking very rapidly, without letting me get a word in edgewise (as if wanting to

add  
“with”

<sup>43</sup> I didn’t realize it at the time, but a three-way meeting wouldn’t have worked at this point, because he and Fritz had already surreptitiously met that day (Thursday), and they’d worked out an agreement to move me off the Wahoo project, as discussed below.

<sup>44</sup> Had I had the chance, my suggestion would have been to fire Fritz, for unforgiveable behavior.

prevent me from saying anything), speaking words to the effect that, “OK,-here’s-what-we’re-going-to-do,-blah,-blah,-blah-...” — it was quite clear Dan didn’t want to hear my input on this matter. Dan said he had two potential alternative proposals in mind. The first proposal (the one he preferred, he said) was for me to switch job assignments with Sujatha Mizer, another member of Dan’s Performance Architecture Group. I like Sujatha a lot, but she is very junior compared to me, and her job assignments (though I knew little-to-nothing about them) were very much beneath my appropriate level of responsibility. I.e., I viewed this proposal to be very much in the nature of a demotion. So I asked Dan what his second proposal was. He twisted his face into a sort-of strange grimace, and said he preferred not to tell me what the second alternative was, and that he “*really*” (emphasis his) thought we should implement the first proposal. He said nothing further about the second proposal, but his behavior led me to believe that his second proposal was probably to “fire Walt”. So I said, “Well then, I guess I’ll choose the first proposal.” Thereupon, he said I could leave his office, and he’d call in Sujatha for a discussion with her, and then he’d call the three of us together later in the day to work out the transition plan. Before leaving Dan’s office, I informed him I was interested in following up “how I should be handling the Fritz thing” with whatever “HR process” was available to me, and he helpfully responded by referring me to Kelli-ann McCabe (as the “legacy Netezza” HR Vice-President), with Diane Adams (as the going-forward IBM HR contact) as back-up choice. I then told Dan: “Do you remember that earlier I told you I thought Fritz was a bully, but that I hesitated to call him a ‘liar’, because to me that word implies an ‘intent to deceive’?” Dan indicated he did remember, and I continued: “Well, I still claim he’s a bully, but based on the ‘yelling incident’, I now see no alternative to also claiming he’s a ‘liar’, i.e., I think he’s consciously promulgating falsehoods with an intent to deceive.” Whereupon, Dan unhesitatingly leaned forward (noticeably deliberately), and looked me squarely in the eyes (this was obviously attention-getting, because he’d been glancing nonchalantly at his monitor, as if scanning email headers, and only partially paying attention to me), and said with exaggerated articulation: “Walt, I need to tell you Fritz comes in to me and says the same things about you.” Obviously, this was yet another major shock to me, because neither Fritz (nor anyone else at Netezza, nor anywhere I have ever worked) had never intimated any such thing to me. I understood this (Fritz calling me “bully” and “liar”) to be Dan’s reason for deciding Fritz and I couldn’t continue working together (even though Dan didn’t explicitly state this was the reason for his decision [he stated no reason at all]). On my way out of Dan’s office, I told him I would indeed follow up with Kelli-ann, adding “perhaps today”. Thus, Dan knew even at this very early stage that I would certainly be pushing the whole Fritz problem through HR channels, and he acted completely sympathetic and supportive, even friendly, about it (as was proper — it was only later that he turned adversarial towards me, see below).

“Mizar”,  
not  
“Mizer”

Dan told me  
not to talk to  
Sujatha  
about Fritz:  
New  
Complaint,  
Add. III, p. 9

OpenOffice  
hyphenation  
bug

see also Add. II, p. 5 and Add. IV, p. 5

Upon leaving Dan’s office (and before the three-way transition meeting with myself, Dan and Sujatha, see below), I composed and mailed the first email included in Appendix I.b; Fritz soon thereafter responded with the second email in Appendix I.b (both these emails contained words of apology).<sup>45</sup>

I then also composed and sent an email to Kelli-ann (including the first two of the three emails of Appendix I, which were the only two published at the time I wrote to Kelli-ann),

45· While both contain words of apology, Fritz’s rings hollow, because he carried through with telling Dan I was a “bully and liar”, with the result that I was demoted, see below. In his email, note that Fritz *self-admits* both provoking the incident, and “raising his voice” (his euphemism for yelling). He does not contest anything in the preceding two emails in the chain, which I composed, and which contain the (correct/truthful) account of the incident and its surrounding context.



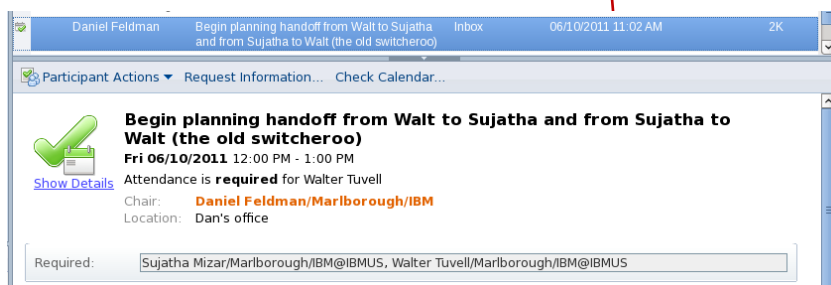
thereby initiating the HR process, as I'd told Dan I would (Appendix J). Kelli-ann responded by setting up a meeting with me for the upcoming Monday, June 13. see Add. V, p. 35

At noon on Friday, June 10, the aforementioned "transition meeting" was held in Dan's office. The substance of the meeting was for Sujatha and I to give a brief outline of the status of our respective job assignments, with one-on-one meetings between Sujatha and I to take place thereafter as occasions presented themselves. For my part in offering an outline of my work, I submitted (paper copies of) Fritz's "upcoming performance tests" (the first email in Appendix H). For Sujatha's part, she offered (paper copies of) her document included in Appendix K. That was really the gist of the meeting. In particular, there was no discussion at all about modifying one-another's projects during the transition — this was also emphasized by the title of Dan's Lotus Notes invitation to the meeting, which was entitled, "Begin planning handoff from Walt to Sujatha and from Sujatha to Walt (the old switcheroo)". We parted (then continued with some other information-exchange, both that day and subsequently).

On Saturday, June 11, I emailed a goodbye letter to my colleagues in Cambridge, entitled "Hail and Farewell". This email, and their replies, is presented in Appendix L. (I also forwarded these to Kelli-ann on Tuesday, June 14.) As can be seen there, all my colleagues denied calling me "liar" and "bully". I believe them, because they have no motivation to lie, and I'd always gotten along well with them. Therefore there can be no doubt that it was Fritz himself who called me "liar" and "bully" (unless, that is, Dan made it up himself). It had been one of my motivations in writing the goodbye letter to determine whether it was Fritz himself had decided to call me "liar and bully", or whether he was just relaying a message from somebody else in the Cambridge office. It is my determination, based on these responses to my goodbye letter, that it was the former.

On Sunday, June 12, I sent Dan my weekly report for the week (Appendix A.dd). I also forwarded a copy of that weekly report to Kelli-ann (Appendix J), in preparation for our meeting the next day.<sup>46</sup>

Also, on Sunday, June 12, I drove to the Cambridge office, and gathered up my belongings, in preparation for reporting for work at the Marlboro office thereafter, beginning on Monday, June 13 (as indeed has happened from that day forward).<sup>47</sup>



k 46. Later, on the morning of Monday, June 13, I sent Kelli-ann a "partial explanation" for some of the wording in the Weekly Report. Appendix J.

47. My workspace in Marlboro was also in an open "bullpen" (much larger than the one in Cambridge), in a temporary/"hotelling" area known as the "Cambridge ghetto". "Everyone" noticed when I started working there regularly, and commented on it (e.g., "Oh, my watch must be slow, it must be Wednesday"), or asked me what was happening, since my normal practice was to work only one day per week, Wednesday, in Marlboro, the other four days being in Cambridge. In this manner, Dan had "put me on display ('in the stocks')", as somebody who'd been "exiled to Siberia". To avoid that impression, I should have been given a "private" cubicle, away from the gaze of the curious.



## APPENDICES — Part I

The following Appendices include many original documents (especially emails), rendered as transcriptions (rather than screen-shot reproductions from Lotus Notes, for example). Therefore some hand-editing has been necessary. Considerations of readability have dictated some insignificant changes (such as omitting *pro forma* headers/footers, and long-form email suffixes such as “/Marlborough/IBM@IBMUS”), but the originals are uncorrected as to misspellings, misstatements, typographical niceties (such as choice of font, straight-quotes vs. curly quote-marks, hyphen vs. en-dash vs. em-dash, etc.), and so forth. The intent of such hand-editing is to enhance understandability, only. In no case has any substantive or misleading change been knowingly made, of course (though some may have accidentally/inadvertently crept in; for example, there may be some “auto-correct bugs,” and some reformatting has been necessary because of the “ASCII unfriendliness” of the Notes platform, e.g., some dates had to be typed in manually).<sup>48</sup> Some editorial commentary has been necessary in a limited number of places to maintain the narrative flow; where confusion might otherwise result, such editorial commentary is indicated by the following special/reserved notation:

►*This is an editorial comment.*◄

### A Weekly Reports

►*This Appendix exhibits all the weekly reports I have generated for Dan to date. These contents come from a text file I maintain on my laptop for my own records, before copying them into email for transmission to Dan. The differences between these and the actual emails I send to Dan are few, if any (though I may have occasionally made a change in the email and inadvertently neglected to copy it back into my text file).*◄

#### A.a Week Of Nov 8-14 (2010)

Finished general G&A stuff (direct deposit, benefits, etc.), and getting set up at both Marl and Camb offices (getting card-keys, setting up printers, etc.).

Continuing/finishing general training. I found the set of videos from Dec 2009 extremely helpful (I'm now finished with them, they should be required reading for all new-hires, and refreshed every few years, say when new/interesting products come out).

Starting to focus on Wahoo-related topics (such as zonemaps and storage layer will change), prior even to learn more about perf tools (such as PerfBar and ODS, those will come later). This is a prioritization exercise, and it seems reasonable to do it this way, as it gets me to start thinking “deeply” about something, instead of skimming over the surface (both have their advantages/disadvantages, but I'm already getting some skimming with the videos).

Next week, planning to understand Wahoo zonemap design to the point where I can make some intelligent comments about the new re-design (this involves some dev-level research into things like computational geometry, though I'll have to time-limit that). Then, start

---

<sup>48</sup> In any case, documents (such as emails) that exist in document-retention databases can/should be accessed in their original form if necessary, given proper authorization. I would welcome that.

touching the Wahoo box in Camb, with an eye toward starting to work with Rich Title's perf harness for Wahoo storage layer.

## **A.b Weeks Of Nov 17-21, Nov 24-28 (2010)**

(First week: short week, 4 days, IBM First-Day Mtg.)

(Second week: short week, 3 days, Thanksgiving.)

(I missed submitting report for the first week above, sorry.)

Per previous plan (with suggestion/help from DanF), I submitted a review of the Wahoo ZoneMap design document to the authors (Devesh, DanD). This was a multi-day effort because it was unfamiliar territory for me, but it was a very useful exercise: (i) helped introduce myself/perf to the Wahoo team; (ii) took me to a deep part of the system (as opposed to the higher-level overview of the Dec. 08 videos); (iii) gave me an appreciation of one of the "four key performance accelerators" of NPS (FPGA, Query Optimizer, ZoneMap, FlowComm, according to Doug Johnson).

Following that, spent a couple of days being tutored (partly with Joseph) in more perf-specific areas by Felix and Larry (they are to be commended for donating their time, and doing a good job).

The second week above, I started a default regime of spending 4 days/week in Camb, and Wed. in Marl. Spent my time in Camb learning how to do dev build, and getting it to run on wahoobox. The latter hadn't yet successfully been done by the Wahoo team, and Rich Title beat me to it (though it was a cooperation, not a competition, and Rich was highly helpful to me and another new person in Camb, Steve Lubars).

I also found out from Paul Smith that a long-standing objective of his has been to support "all" (reasonable) flavors of Linux for the engineering team, but he'd been limited to RedHat and Ubuntu because that's what most engineers use. Since I use Fedora, he now has some ideas about restructuring parts of the wiki, with my help filling in some Fedora details. I anticipate this will take only a small part of my time, and if it turns out to be burdensome I'll re-evaluate my participation.

Next week, am planning to be part of Fritz's next "iteration" (or "sprint", in scrum-speak), and work with Rich to start understanding the storage-layer perf tools he's put in place.

## **A.c Week Of Nov 29 - Dec 3 (2010)**

►*It appears I missed sending a weekly report for this week, it's missing from my records.*◄

## **A.d Week Of Dec 6-12 (2010)**

Got NPS to build/run, in Wahoo VSPU mode, on both Ubuntu and Fedora. This involved updating some of the instructions published by Paul Smith on the wiki for Ubuntu, and reinterpreting them for Fedora.

For the record, I'd been told: (i) by corporate IT that nobody runs Linux natively on their laptop; and (ii) by several people that engineering runs only Ubuntu, though some people had

tried using Fedora and abandoned that effort. So I decided to plunge forward and find out what the problems were, on the theory that (i) doing Linux on laptop would help me understand the Netezza environment, and (ii) bringing up NPS (VSPU) on Fedora would help me understand NPS better (at least the build environment, but I also needed Wahoo VSPU to play with PerfBar & friends until a real Wahoo box exists). In the process, I've been working with Paul to update the wiki, and while I've sensed (perhaps mistakenly) some friction doing that I'm hopeful we can work it out.

During the week attended many meetings of one sort or another. Especially useful were Felix's two long brain-dump seminars on PerfBar (though Joseph & I'd seen much of it informally before, the reinforcement was great). Certainly, much mtg-going is important for getting embedded into Netezza, but it also contributes to my frustration about "showing I'm doing something" (I'm feeling this is reaching a kindling point with Fritz, esp. w.r.t his Scruminess).

Towards the end of the week, wrote a csvBeautify.py program to make PerfBar reports readable without Excel (or OpenOffice Calc). Again, this might seem like a minor thing, but it sure helps the on-boarding process.

Next week: Start hacking with PerfBar & friends in my fake/VSPU environment, with the goal of figuring out how it can help us with Wahoo. (It's because of a desire to get this part of the project going that I didn't want to accept Dan's suggestion of running some SPU-elimination tests, as discussed during Gordon't mtg on Fri.)

## A.e Week Of Dec 13-19 (2010)

As planned, I worked on getting PerfBar running on my laptop/desktop VSPU Wahoo environment. The goal is to learn enough about PerfBar so that I can pick-and-choose-and-modify some tests for a very specific purpose: a Wahoo deliverable to the program committee in early Jan, to "present at least one comparative recognized performance number". This means using some PerfBar test (recalling Wahoo doesn't yet have enough functionality to run the whole PerfBar) so that we can exhibit that it's faster than a comparable disk-based NPS (recalling again that a Wahoo box doesn't yet exist, and the comparable disk-based NPS will have to be hobbled for comparability).

That plan is working, IMHO. By week's end I was able to run the TPCDS first/initialization phase of PerfBar (i.e., the 401\_create\_tpcds\_db.py script), and at least one PerfBar test from the TPCDS second/workload phase (i.e., a SQL script excised from the runBar.py script). Can't run the whole of runBar.py yet, or even run that excised SQL script from inside PerfBar, or the other (non-TPCDS) suites. But that's OK: all we ultimately need is "one comparative recognized perf number", by early Jan, so all-in-all this is sufficient progress at this point (and, BTW, another "first", in the sense that it hasn't been done before, though others had tried).

This week felt better, w.r.t. "doing things" as opposed spinning in meetings. There's nothing wrong with meetings, but when they interfere with higher priorities, there's a problem. The one exception was on Thur, when Fritz held a 4-hour (!) Scrum planning mtg. After it, I was anxious plant myself in front of a keyboard (Dan, that's when you saw me bolting out to grab a sandwich). Actually, the whole team wanted to see some changes made on the Scrum front (though it has its good points), and Fritz has been allowing some changes on an experimental basis. We'll see.

## A.f Week Of Dec 20-26 (2010)

(Short week, Christmas.)

Got essentially all of PerfBar working on laptop, on Wahoo VSPU. This means got over the final humps that were left over from last week, i.e., now hitting no systematic problems, just sporadic ones. (That is, porting perfbar is complete, but there are still some bugs or known problems in PerfBar and/or Wahoo.) Whether or not the porting changes, or a subset thereof, will be folded back into PerfBar remains to be seen. I have a full set of notes, and will consult with Larry/Felix, but that is a secondary task at this point. For now, continue learning fine point of PerfBar for purposes of Wahoo, see next paragraph.

Main task remains getting some perf numbers for Wahoo program mtg in early Jan. Wahoo is just now making its trek to prototype machine, so the whole push will be a parallel and essentially full-time effort for the next couple of weeks. To that end, gave a chalk-talk to Wahoo team on what PerfBar is all about, and to chart a course with the team towards what we need for program mtg.

Finished setting up desktop in Cambridge. Also, got new IBM laptop and started figuring it out, very nice hardware, but not happy with preliminary software situation. IBM has a semi-supported Fedora ("Open Client") install, but it's very restrictive, for example, only seem to allow IBM-massaged downloads/installs, and doesn't allow root access to the user (and, encrypts the disk so the user can't manually remove passwords and bypass that restriction). But there may be loopholes (legal ones, specified in the IBM small-print, such as whether the user stores user-data, as opposed to mere company proprietary data) to get around this, so I'll keep looking. Hoping for happy outcome. One curious factoid is that IBM seems to prefer Fedora over Ubuntu, thus pushing Netezza people in that direction, which is what I'd already been trying to do too.

Sidelight: When Sujatha copied Mark Odell's IBM Fedora Open Client for me, she did it by copying files, not the ISO, so I couldn't use it. But for Christmas, my 17-yr old daughter received a Garmin GPS for her driving (her first-choice Xmas present), and wanted to test it out, so I sent her on a Trek out to Marlboro to retrieve the DVD. Made both our days.

## A.g Week Of Dec 27, 2010 - Jan 2, 2011

(Short week, New Years.)

Finished up (well, as far as anything is ever finished up) the PerfBar tasks needed for Wahoo's conjectured program mtg on Jan 7 (though it seems the actual mtg will be a couple of weeks later). That is, got a sufficiently large subset of TPS-DS to run on both Wahoo and Skimmer for comparison purposes. The "at least one comparison test" will therefore be "the subset of TPC-DS that runs reasonably on Wahoo and Skimmer". That subset consists of > 50 tests, larger than thought previously, because Wahoo now only lies (just returning true, or the moral equivalent) on operations it had previously crashed on (esp. GENERATE STATISTICS and DROP TABLE, which PerfBar does a lot). Reasonableness means excluding tests that Wahoo doesn't support (DELETE, UPDATE), or that just plain have weird problems (test #74 hangs on Skimmer, probably a PerfBar bug, not necessary to figure out this problem just now).

Note that Wahoo software still doesn't run on Wahoo hardware (which doesn't exist yet), or even on Wahoo prototype (which still doesn't exist yet either, some combination of OS not having drivers for Violin and NPS not getting through its PXE-like boot-up sequence). This means the desired perf comparison numbers surely won't exist in time for Jan 7, and are even in jeopardy for a couple of weeks after that. A couple of weeks ago at Fritz's scrum planning mtg, I had volunteered to help with the bring-up effort, but was rejected, because he thought I was already a bottleneck, but that's turned out not to be the case. This comment is in no way intended to imply I could have done better than others, only that I continue trying to be a good citizen (indeed, had I worked on the bring-up, I'd be behind on PerfBar). It also shows that I'd already tried to get more involved in the design/implementation of the prototype. At the time I volunteered, I recognized not only that the bring-up was the bottleneck, but also that I needed to understand it better for perf/comparison purposes (just as Dan has also urged upon me). [Quietly: Fritz is now waffling that the effort to bring up Wahoo on the prototype is even well-advised, think it might have been better all along to just press directly forward towards real hardware, or some other prototype design; unclear what direction this thinking will take, should be become clearer next week.]

Towards the end of the week, started migrating from Dell laptop to Thinkpad in earnest. Making good progress on the parts I can control, that is, haven't gotten into the IBM-specific stuff yet. I'm making good re-use of the earlier Fedora notes I'd made for the Dell, fixing bugs/typos (these same notes have also been useful to James Percent in Camb, and to Joseph Shkolnik who hadn't used VMs previously). These notes are still the best intro to building a Netezza-ready Linux box, IMHO, but now that IBM is in the picture Paul Smith and others are re-visiting the whole area it's unclear what role my notes will play (had already discussed all this with Paul pre-IBM, and will continue to follow it up, backburner).

Happy New Year!

## A.h Week Of Jan 2-9

(Short week, New Years.)

IBM and ThinkPad migration. Almost done at this point. Hit some dead-ends, like many in Dev, but I hit them earlier and I like my solutions better than theirs at this point, so I seem to be ahead of the curve (amongst those seeking non-standard solutions). Though jury is still out, as to whether my solutions will continue to hold up (I like my chances).

Spent a number of hours with Joseph, getting him up to speed on Linux and VM technology. I like to think we're succeeding

Had a talk with Fritz about direction to move forward with Wahoo Perf. We're of a mind. Am also inching closer to Dan's worthy goal of characterizing Wahoo in an arguably intelligible way, vis-a-vis similar systems, a major goal of mine for upcoming week.

Doing a fair amount of side-reading on data-warehousing & friends (OLAP, DSS), and how it differs from classical DB (OLTP). Actually have been doing this for awhile, and will continue (I find the time during odd hours, such as first part of morning work-outs when stretching, and at breakfast).

Toward end of week got a good start on "mugshotting" (a.k.a "fingerprinting") NPS systems, a necessity for almost any perf study one can think of.

Sidelight: Got a new set of nice kitchen knives for Xmas, so fingers have gotten cut, and am typing this with a bandage on my finger, though not so bad as my son, who took 3 stitches.

## A.i Week Of Jan 9-16

Ran several Wahoo (& comparative) PerfBars, in preparation for pgm mtg on Wed Jan 19. There are some problems for Wahoo: (i) failures (negative number) on some PerfBar tests; (ii) there has been perf regression; (iii) hangs in mid-PerfBar; (iv) untimely time-off for the principal (storage-layer) players (vaca, sickness, IBM training, moving-day). Will be a photo-finish for pgm mtg.

Enhanced csvBeautify.py to produce comparative totals. It now adds up cols, using only the rows that have all-pass (positive numbers) for the whole rows, where each col represents a run you want to compare. I don't know how easy that is to do on the usual spreadsheets (Excel, OOo), because I haven't tried.

Getting near finish line with nzMugshot.sh. Got good feedback, incorporating it (only Gordon's conception was far off the norm). Should finish this week (background activity, no hurry).

Working on Wahoo commensurability, starting from Dan's spreadsheet. Yes, I think it's important, just not urgent. Plan to make serious progress this week.

## A.j Week Of Jan 16-23

(Short week, "Succeeding@IBM".)

Ran some PerfBars against WahooProto and Skimmer, in preparation for pgm mtg. Numbers were disappointing, but seem to be helping the dev team. I.e., it seems PerfBar is working OK, and the problems are believed to be in NPS/Wahoo, since the numbers are being replicated by dev-tests. Not obvious, and some divergence of opinion. as to what's causing the perf problems, starting to whittle down the candidates.

Need to talk go Fritz/Dan about best ways to contribute next. Fritz has suggested with a menu of additional tests to be done, and he's probably right about those being the best next steps, though need to confirm.

Continuing to work on the Worksheet Model initiated by Dan. It's now reaching an interesting point, hence it's a good time to invite Fritz's ideas. Additionally, cobbling together a visual (Visio, Inkscape) that brings things together better for me than I've seen elsewhere.

Still got some ThinkPad setting-up to mess with, such as getting VPN to work (non-trivial, given that I'm all-Linux), and getting more memory.

## A.k Week Of Jan 23-30

(Short week, 2 half-days of IBM Open-Source and Patent training.)

Continuing work on multiple projects: csvBeautify.py, nzMugshot.sh, "spreadsheet model" (including a new visual, helpful to me and I hope ultimately for others too). These things are coming out nicely, but are going slow, not because they're super-hard or I'm super-dumb, but



because I'm timeslicing myself to death. Larry pointed me to DougJ's previous work on a spreadsheet model, and I've contacted him to pick his brain (if he's going to be given a new contract).

Next sprint/iteration planning for Fritz. Another half-day gone.

Studied and started a discussion about our use of "geometric mean", and the possibility of replacing it with a "compare-to-Skimmer" metric. This has lots of theoretical and practical advantages, cf. the MIPS metric (= "compare to Vax 11-780", see Wikipedia article on MIPS).

Upon DanF's instigation, Michael Sporer had added me to the review list for Bruno DiPlacido's RTR hashing study, and I had some thoughts about it, which I shared with Garth. Garth liked it well enough to suggest I work with Bruno, so I wrote a script to generate the sample distributions for Bruno to test with. It uses a flexible user-written specification to create arbitrary weighted random distributions (as opposed to uniformly random distribution) of integers (but could be adapted to other datatypes). Such input are the critical ingredient to the RTR hashing work of Bruno.

Questions surrounding usage of OpenOffice, and other Open Source Software, still in process. Researched "Official IBM Blue Color". Sorry I did.

## **A.l Week Of Jan 30 - Feb 6**

(Slightly short week, was a bit sickish a couple of days.)

Finished migration from Dell laptop to ThinkPad, and from Outlook/Exchange to Notes/Domino. I like Notes better than Outlook, except that perf seems to be much worse, which is almost a deal killer.

Also finished migration from Dell workstation to Thinkpad. No need to maintain 2 machines when 1 will do (just a maintenance headache). Tried to loan both Dells to Camb folks (they're short of hardware for Wahoo), but they weren't good enough (they want 4 cores, not just 2 [ignoring hyperthreading]).

Wiped Dell laptop clean, will return it to Marl IT next week.

Worked on Perf Report doc template, at Dan's instigation, with other Perf Group members.

Filled out self-eval, and turned it in.

I think I may be have some dissatisfaction with how Perf Group is doing some of its reporting. Geometric mean may be only the edge of the wedge. Am thinking more about this.

Notably, didn't do much for Wahoo Perf. That's because Wahoo's problems are now at dev/impl/debug level, where I'm not currently qualified to help out. Without placing fault, this is generally a point of concern: am I being used, and/or contributing, the best way I can? Ongoing dialog with Dan about this.

## **A.m Week Of Feb 6-13**

When I was running PerfBar TPC-DS for Wahoo a few week ago, it was throwing some sporadic errors (in the sense of reporting negative elapsed times). When I tried running the test

on NPS 6.0 on Skimmer, I also got some errors (some different, some same). At the time I chalked this up to "technical details that probably wouldn't turn out to be very interesting" (whether the bug lay in PerfBar, Wahoo, or myself), so I delayed investigating. This week I decided to take a closer look. Working with Felix, it turns out that PerfBar itself is at fault. The problem is indeed a "technical detail", a bug (err, "unadvertised feature") in PerfBar that doesn't really hurt what we need for Wahoo (namely, PerfBar only "really" works with ballooned DBs, but Wahoo can't do ballooning yet, because that requires UPDATE, which Wahoo doesn't support yet).

Also started to take a closer look at some of the other burrs that have been under my saddle for awhile about our perf stuff, esp. (i) detailed understanding of TPC-DS, and (ii) how to present/report our results for the best benefit of others, both numerically and graphically. I solved (ii) to my satisfaction (haven't communicated it to anybody yet). Discovered Felix is already embarking upon the same (ii) journey, so I'll collaborate with him (unobtrusively). Item (i) remains in progress.

Spent a half-day studying up for Ashish's Queueing Model review mtg, came up with a solution I liked. Then attended his mgt, and he came up with exactly the same solution. Dan also agreed. Good sanity check for everybody.

More work on the Perf Group Doc Template (semi-background project). Nearing consensus, it's looking quite nice.

There's contention for the Skimmer I've been hogging, so I may need to abandon it, for awhile at least. Since I had previously gotten NPS/VSPU to run on my Dell laptop, I'm looking into doing that trick again on the ThinkPad (since I don't really need a real Netezza machine to develop some of my perf ideas). By week's end hadn't made much progress, due to time-slicing/context-switching.

Still unresolved is how exactly I'll be proceeding w.r.t. Wahoo. Late-breaking news at end-of-week is that Dan Dietterich has finally made the break-through on understanding where the cycles are disappearing to in Wahoo vs. Violin, so hopefully I'll be able to enter that arena again soon. Fritz also said he has some ideas in mind. Will find out early next week, when next "iteration" planning happens.

Measured the amount of data returned to Host during TPC-DS test run. This was needed to convince ourselves that the anemic WahooProto Comms fabric isn't causing slow PerfBar numbers. It isn't.

Joined in the "secure disk wiping" email free-for-all, by distributing the most recent paper I know about on the subject. It contradicted what Brian Maly had earlier posted, hope he doesn't hate me. (Actually, I privately asked Brian about what he wrote, but his answer was dismissive, so I was "forced" to send my post anyway.)

## A.n Week Of Feb 13-20

Perf Report doc template: finished it and posted on wiki. Fine.

Doc-reading & review mtgs: Hermes; Hash Study; FPGA Perf Test; various QA specs that cross my desk (manh of these just got a cursory glance, as I continue building up general Netezza context).

Iteration 13 planning in Camb (Tue). Now that the weird Violin/vRAID perf issue has been resolved, it's time to do some more regular perf testing. It has been decided to start with lower-level storage-layer testing, using DanD's PMtest (Partition Mgr), as opposed to the abortive PerfBar testing that was tried earlier (will get back to that later, when Wahoo is a little more mature). By end of week, got some PMtesting done. It's painful because Wahoo is still unstable, so the cycle is: (i) a few min of data collection; (ii) 20 min rebooting Wahoo and creating a test DB; (iii) lather, rinse, repeat, except that this can't simply all be put in a script because of unpredictable Wahoo hangs.

Nearly a whole day (Thur) doing stuff with Clark French. This was unexpected, but a lot of fun, and I hope it'll continue being a lot of fun continuing it for the next few weeks, scheduling appropriately with Wahoo duties.

## A.o Week Of Feb 20-27

(Short week, President's Day.)

Worked over WHOLE weekend and Prez Day (family was out of town), on Wahoo/Violin perf measurements & report. Fixed broken network/UPS in Camb in the process. Got draft-of-draft out for review. This work is already showing up some anomalies nobody suspected, and hence will certainly lead to a longer-term project than expected (that's Fritz's problem, not DanF's).

Filled out Netezza Quality of Life survey (or some such), put out by IBM Integration Central. In their free-form answer sections at the end, I wasn't shy about suggesting "Let Netezza be Netezza".

Worked some on project Glass, with Clark French. Larry help by providing me with the DDL (I don't know how, later discovered it's done by the usual way of asking the DBMS, which I don't have access to). HOWEVER:

Took the IBM (required) Info Sec course. Didn't learn anything, except that IBM pretends to be serious about security. But ran into a roadblock trying to follow the rules so I can work with Clark. If this doesn't get fixed first thing Monday morning, I declare it a blocker, eligible for escalation.

## A.p Week Of Feb 27 - Mar 6

Whole week (apart for normal/predictable/ongoing/recurring activities, such as spare-time reading, the usual Wed in Marl devoted largely to mtgs, and the recurring Sprint stuff with Fritz) spent on 2 activities:

1. PMtest'ing for WahooProto/Violin. Now have the full complement of raw data (25,600 runs of PMtest, in 40 clumps of 640 runs/clump). Have graphed & analyzed these in barplots. What remains to do: dream up and generate additional reporting vehicles, some graphical and some numerical; write the whole thing up. The writing-up involves some give-and-take with DanF, and presumably other procedural steps TBD.
2. Project Glass. Finally got secure comms squared away with Clark French. What we really want is secure email, but what we've got is secure FTP dropbox, with email notifications of dropbox activity, augmented by higher-bandwidth phone conversations. This is working, but

it's egregiously painful. It's a procedural nightmare, including the issue rich-text emails I have with people internally that has to be converted to FTP-able files. Just a part of doing business in the New Regime, but frustrating nonetheless. In any case, Glass is now up and running. I don't quite follow some of the things Clark is doing, but I'm following along until I get a better grasp, and in any case being as responsive as I can to anything he wants in the meantime.

## A.q Week Of Mar 6-13

Repetition of last week: whole week (apart for normal/predictable/ongoing/recurring activities, such as spare-time reading, the usual Wed in Marl devoted largely to mtgs, and the recurring Sprint stuff with Fritz) spent on 2 activities:

PMtest'ing: Circulated 2nd draft, got only a few comments back, but still need to at least (i) generate more graphics (cross-sections through other dimensions, per Garth's suggestion, but it's a good idea no matter whose idea it was), and (ii) generate some numerics (I've got some ideas, but if I don't have enough time I may cancel the whole idea, on the theory that most people will be happy enough with the visuals).

Glass: In the midst of this project at this point. Doing my best to keep Clark's queue full, and succeeding so far, but he's fully engaged at this point, and moves pretty fast himself.

Concern: I'm starting to feel like I'm being pulled to do new things on Wahoo, without being given enough time to finish-off PMtest & Glass (esp. w.r.t. things that get added here and there, such as perfectionism on the PMtest doc, and additions to Glass such as CTA2 and more scrounging around in data than initially expected). That's not to say these aren't things that things that shouldn't be done, just that it takes time to do anything, even on the downhill side of projects. If things start feeling untenable, I'll complain louder.

## A.r Week Of Mar 13-20

Continuing with PMtest Report. It's a more major effort than anyone thought. Right now, the task at hand is to find a way to visually represent a 4-dim hypersurface in 5-space. This was Garth's original wish to me weeks ago, and he made the "mistake" of saying he didn't know if it was possible. Being intrigued by the question, it grabbed my attention, but I wasn't able to work sufficiently on it because of Glass. But now, since the Glass duties have been lifted, I actually have found the solution late last night and earlier today (Fri-Sat). Will still take some non-trivial number of hours to implement it, but the back of the problem is broken. The solution involves non-trivial work with doing data-graphics using the R statistics system. After that, still have to figure out what kind of "numerical analysis" to do for PMtest. Sigh.

Tried to do another round of PerfBar/TPC-DS on WahooProto. Couldn't do it earlier, because either: (i) it was futile because WahooProto was known to be uselessly slow (before the "randomized scanlist hack"); (ii) it was over-subscribed due to merging; (iii) it would crash in the process of trying to build the TPC-DS database. Once those problems were fixed, I did a TPC-DS run, but it hung after ~20 queries. But those queries revealed another perf bug (still ~20 times slower than Skimmer).

In the throes of problems (i)-(iii) above, Fritz sent me on a wild goose chase to ASCII-dump a TPC-DS DB from another machine, and hand-load it into WahooProto (as opposed to regular

backup/restore, because WahooProto doesn't support restore yet). I was upset about this, because I was already over-worked trying to finish up PMtest and Glass. As a result, I had to relinquish my participation in Glass. Sad day.

Also presented to DanF my idea about establishing Skimmer as our "unit of (performance) measurement". He agreed with the concept, and I plan to implement it anon. I think this is a Big Deal.

## **A.s Week Of Mar 20-27**

PMtest doc is in very good shape. Only serious remaining task (i.e., other than TBDs) is to add numerical stuff and recommendations. The numerical stuff is expected to be quite small, and I know exactly what I want to do for recommendations, though it'll take a little calculation to get the right answers.

Along these lines, the way I've done multi-dimensional data visualization seems to be novel (at least a cursory investigation hasn't turned up disqualifying prior art), so am toying with the idea of filing a patent application. Not in the company's sweet-spot of competence, but might be fun to pursue it. Due to IBM's well-oiled patent machinery, I anticipate that would take relatively little of my time, if I do so.

In any case, I can't finish up the PMtest doc in one push, because I keep getting side-tracked by that pesky Wahoo project. (That's a joke of course.) There, the team is tackling the Grand Challenge of figuring out what that slippery perf bug is, and they're committed to killing it once and for all. That means my next immediate project will be to the long-delayed engagement of PerfBar'ing Wahoo.

That's good, because it dovetails with 2 project I've long been wanting to tackle: (i) defining the Skimmer Perf Standard; (ii) properly documenting PerfBar. This week, I started doing both of these things, specifically, outlining to DanF my ideas about (i), and starting a document about (ii). Much more to come.

## **A.t Week Of Mar 27 - Apr 3**

Most time this week spent on Wahoo PerfBar testing. Learned that JIT-stats don't matter for TPC-DS or Atomics (and some other components of PerfBar), unless NPS has major bugs in its JIT-stats implementation, as Wahoo does. But most importantly, the perf bugfixes in Wahoo in recent weeks/months has now finally made WahooProto run at parity with Skimmer. This means we can now start doing perf work on Wahoo in earnest (up to now that's been pretty much useless), because WahooProto is supposed to run at 10x Skimmer. (Yikes!)

Finished (!) the PMtest document, in all essential respects. Still have some minor clean-up to do, and need final review/approval of DanF. Those should happen next week. See attachment.

Started working in earnest on the new/improved perf-reporting structure. Have verified its viability, and even gotten some proof-of-concept graphics. Hope to start socializing this stuff very soon, perhaps the upcoming week.

Still have plans to do proper documentation of PerfBar, time permitting. Depends on how involved perf work on Wahoo gets in the short term (in the long term, we know it'll be all-consuming).

## A.u Week Of Apr 3-10

Mostly Wahoo testing. For the time being we've abandoned our overly-ambitious plan of going-for-the-gold with TPC-DS tests, and backing down to Atomics. Atomics tests are simpler, giving us a better chance to break down into easy pieces and analyze those. Which means, I'm getting more into the job of understanding plan files, and the snippets contained in them. There was no need to have to do this before now, because Wahoo wasn't ready to accept it, but now's that time. Biggest progress this week as been configuring systems such that Skimmer and WahPro generate/run the same exact plans/snippets, so that we can compare apples-to-apples (first defeat Skimmer, then the World). This has now been achieved (provisionally) by turning off JIT-stats and 2-phase-planner (because Wahoo is buggy/immature in those areas), and using same number of data-slices on both machines (8). Verified that GENERATE STATISTICS is generating the same statistics on both machines, so we'll be leaving that turned on. There's still some suspicion the 2 system are working differently at low-level storage layers, but don't yet know how to verify/attack that. Will look at how to get some better system monitoring next week, and instrumentation later on. Doing all this on a test case (EXCEPT test) should generalize to other tests.

Published PMtest doc, to worldwide acclaim.

Got exasperated with Joseph's incessant whining about anything and everything, and blew my top. A little. Will try harder to avoid in future.

Got re-involved in clock-synchronization discussion. Recall I'd visited that a couple of weeks ago at Camb/Wahoo, and then Joseph raised it independently at Marl/Perf. He (and apparently Ashish to some extent?) carried the water for awhile, but the report he (they?) submitted was woefully garbled and inadequate, so I picked up the cudgel and finished the job properly. Submitted Jira ticket for Netezza system to provide synchronization at sub-second accuracy (externally via NTP, not internally via NPS), but in the short term wrote a script that does the trick at second accuracy.

Next up (not counting Wahoo): Further work on developing Skimmer Standard (this includes GeoMean issue, and better perf reporting regime). Intend to work on it during interstices in Wahoo tasks. This project is already started, and I know it's "right" and exactly how it's going to end up. But knowing and doing are 2 different things.

After that: I still think we need "proper documentation" of PerfBar, because that way we can get other (non-Perf-Group) people on-board with interpreting PerfBar results. TBD.

## A.v Week Of Apr 10-17

Distributed first draft of Skimmer Standard proposal doc, with follow-up email discussion amongst the reviewers. Grateful for support I'm generally getting (though there are always exceptions). The airing of concerns by Garth and Larry has led to important fodder for discussion, and it'll be folded into the next draft. Most importantly, Skimmer is being exposed as not quite sufficient for unit-of-measurement, largely because of scalability insufficiency.



This doc will remain my 2nd most important task for the next week, because I think it's important to do this properly -- not only the rigorous/numerical mindset of using meaningful one-number-summaries, but also for the potential that the resulting qualitative good visualization/data-graphics have for everyone in Netezza.

Continued work on Wahoo perf, esp. delving into details of a single targeted query, the Atomics EXCEPT query. We (both me and Camb-in-the-large) continue to learn new tricks, such as running same dev branch (dev-wahoo) on both machines for engineering studies, though that's now how we want to "compare Wahoo to Skimmer" (that will use mainline NPS 6.0 on Skimmer, not dev-wahoo branch). Also plumbing the info in low-level logs, such as `SpuPerfLog` and `_vt_disk_log`. This tack (of beating one horse dead before proceeding to the next) has been fruitful, and seems likely to continue for awhile longer (after EXCEPT is beaten to death, move on to the next-lowest low-hanging fruit), before we attempt to generate another roll-up (full run of PerfBar Atomics or TPC-DS, and compare to Skimmer).

Also, aiding some others (Rich, DanD) on how to run PerfBar on their pet projects, as a sort-of smoke-test in anticipation of true QA involvement. Even had a talk with Devesh about Skimmer-Standard/geo-mean, and Fritz mentioned it might be good to give a lunchtime chalk-talk to the whole Camb about it (not sure how serious he was, but if he wants I could do it with zero prep time).

## A.w Week Of Apr 17-24

Whole week spent doing just 2 things:

Final draft of PerfScore doc, including publication.

Doing/reporting WahooProto perf numbers. This is still drilling down in the dark recesses of one query, the Atomics EXCEPT test. We (not just I) are all still learning about this stuff, including exactly what kind of stuff we want to learn about. Progress.

## A.x Week Of Apr 24 - May 1

Whole week spent doing Wahoo perf. Still focusing on Atomics EXCEPT test, in ever-deepening levels of detail.

At beginning of week, this deep-diving culminated with the gathering of sufficient detail (thanks to Devesh's long-dormant monitoring/instrumentation programs), and of sufficiently understandable reporting (thanks to me), that a mtg with Fritz, DanF and me led to an indication that Wahoo was being CPU-bound, not storage-bound -- contrary to conventional Netezza wisdom (and hence a core design principle driving products, such as Wahoo). This was only the single query (EXCEPT), and a perhaps-not-representative one at that, but Fritz freaked out, and started grasping at alternative design strategies, such as changing CPU/storage ratio.

But one test does not a performance report make. Therefore, the remainder of the week was spent sharpening perf reporting tools, so that the EXCEPT-type report can be repeated for other Atomics tests, and then branching out to other Atomics tests, and to TPC-DS. The tool-sharpening activity consists mostly of scriptifying what has been heretofore mostly Emacsifi-

cation. By week's end this scriptification is nearly complete, and will be able to start on gathering/reporting perf numbers next week.

Soon (after Wahoo perf work is sufficiently satisfied): PerfScore reporting tool.

## A.y Week Of May 1-8

(Short week, day off on Mon.)

Spent the whole week polishing up Wahoo Atomics scriptization, and running same. The latter was "under pressure" (i.e., before final touches were put on the scripts), because Fritz wanted some early results -- better are coming, which is why I haven't sent the early ones to DanF yet.

This involved finding & fixing bugs in PerfBar, namely, PerfBar advertises that it works for non-ballooned TPS-DS DB, but it doesn't, and since Wahoo can't balloon yet (because that required UPDATE capability), I had to fix it.

Next week: Finish that. Then start on PerfScore, and/or Wahoo TPC-DS, in some order.

## A.z Week Of May 8-15

Mon-Wed spent finishing up Wahoo perf-report scriptization, and publishing the results for ~20 of the Atomics tests. This task being put on shelf temporarily, due to lack of machine resources (WahooProto is being upgraded to dual-headed box, the 2nd head of which doesn't work because the xSeries can't see the new Violin device -- Brian Maly and Violin staff on-site in Camb looking at it). When system becomes available again, planning to run full set of Atomics tests [checking new numbers are comparable to old numbers, given the change in headedness], cleaning up remaining buglets entertaining new featurelets in scriptization, analyzing, moving on to TPC-DS tests, and all that. I.e., full-swing Wahoo perf work.

Thu-Fri spent on writing PerfScore graphing/reporting tool (gear-shifting requiring the usual context-switch overhead). Using R wrapped inside Python, which is fugly at coding level but effective (users won't see the implementation), and should even have the serendipitous side-effect of automatically making the tool usable in any environ that has a Python interpreter, such as Window. Should be done next week (recalling the graphing part has been written heretofore, not a user-friendly script, not the stats-reporting part, nor to mention late-breaking reqts from Michael Sporer, also next week is shortened by me taking off Thu-Fri). Will get some other people (e.g., Joseph) using it, etc.

## A.aa Week Of May 16-22

Spent the whole week, really, doing PerfScore. This includes writing/debugging the tool itself (perfScore.py), with a test tool to exercise it, writing v1.1 of the PerfScore.pdf doc, and packaging it up with an installer and instructions. Sent it out for beta-testing, in the hope that somebody other than me will actually try it out. Once that is done will announce it for general distribution, hopefully next week.

Oh yes, there was a new kerfuffle about Fritz, who apparently claims I now need to have an expertise in mind-reading. Have escalated the issue to DanF. I won't be surprised if Fritz

succeeds in getting it swept under the rug as a "mere communication issue". I speak/understand English, I don't know what language he's using. Call me a cynic.

Next week: Back to Wahoo perf testing, now that there's some machine time available. I think/hope. But it's likely to be a short week for me, with deck refurbishing and all.

## A.bb Week Of May 22-29

Aim to get a new comparison of Wahoo vs. Skimmer on TPC-DS. Gathered the required data from Skimmer (1/100/1000 GB) without incident. Wahoo is another story. Various Wahoo software breakage caused us to downgrade from the new 2-headed virtual Wahoo1/Wahoo2 setup back to physical WahooBox (this was our first experience on virtual, so we wanted to eliminate that variable), before we discovered the main problem was a bug in parallel-load code. Proceeding with single-thread load, we got 1G numbers. But loading 100G took much longer than expected (what should have taken ~1.25 hrs expanded to >36 hrs). By week's end (Fri), the 100G was loaded, but testing not yet initiated.

Nearing completion of PerfScore tool. It's code- and doc-complete, now doing testing & bug-fixing. Devesh is helping with that, he's done a good job and found some bugs (on a system different from my Fedora, of course): my bug in the installer (wasn't building R shared-lib); some kind of functionality bug (R graph generator isn't finding the non-ASCII Unicode chars in Deja Vu for some reason, not yet resolved).

While functionality-testing PerfScore, found a bug in R (floating-point round-off error when running stats on vectors whose components are all identical, esp. single-component vectors), will report it.

For gathering system stats (host & blade), looked at nmon (DanF's suggestion), and modified it for Wahoo's needs (0.1-sec sampling and timestamps), on Fedora. Contacted the nmon owner (Nigel Griffiths, IBM UK), but he hasn't responded yet (we don't actually need him at this point, but he should incorporate our reqt into his code). Still need to port nmon to Wahoo host & blades, and parse CSV files (perhaps using exsistant open software tools, perhaps will write our own). Modulo that remaining work, it looks ideal.

Talked/harangued with DanF and Fritz about "perceptions of work". Hopefully got that cleared up, at least to the extent of how to proceed in future.

## A.cc Week Of May 29 - Jun 5

(Short week: Mem day, and some off-time for medical purposes [was supposed to work on deck refurbishing, but carpenter go sick].)

Was supposed to get TPC-DS@100 numbers on Wahoo by now, and spent most of the week trying to do that, but it's been impossible despite many tries, due to broken dev-wahoo stream (both now and attempted time-travel-backs [Accurev sucks]). By week's end, decision was made to try fixing dev-wahoo, we'll see how that goes.

Continued working on 0.1-sec version of nmon, now called nzmon to avoid confusion. It now contains some additional innovations for time-keeping, esp. "micro-adjustments" to ensure 10 snapshots per second despite the tendency of clocks to drift. Also started looking at generat-

ing reports from nmon capture files, esp. via RRDTool (as opposed to Excel), but with limited success (the stuff doesn't quite work as advertised, not surprising for open-source).

Updated Fritz and DanF far too often, "as if" this weren't a typical week in Wahoo-land, but it was.

## A.dd Week Of Jun 5-12

After 2.5 weeks of Wahoo total noncompliance of various styles, the machine finally came to life, i.e., became usable on Wed for perf testing. I ran an overnight test Tue, and Wed morn saw it had succeeded. Running PerfScore tool on it, I accidentally key-bounced Wahoo 100G vs. Skimmer 1000G upon CLI entry (auto-complete), which resulted in a false indication that Wahoo was >4x faster than Skimmer, which I (erroneously) reported. There was singing in the streets of Cambridge. But I discovered my error late Wed night at home, and immediately reported that the correct reading was that Wahoo is 1.7x slower than Skimmer. A major part of the bug was that I "wanted" Wahoo to be fast, and deceived myself, instead of coolly sanity-checking. As a double-check, I carefully re-ran the 100G test on Thur morn (coming into the office at 5:15 AM, skipping my normal morning workout to do so), concurrently on the 2 machines, visually observing their progress. This confirmed the 1.7x slower number as definitive.

Late on Wed afternoon (before my error, above, was known), Fritz published to me and Steve a list of 5 "upcoming performance tests", which I completed by Fri (and I assume Steve completed his part too). Actually some of the things Fritz suggested were worthless, as "must" have been obvious to Fritz, but it's his manner to arbitrarily assign scut work to me (seemingly due to neuroses of his own, as has become increasingly clear to me). These were things like running debug builds where there was no reason to, and collecting stats via PerfBar when those had already been collected by WaltBar runs. So I gently pointed out their worthlessness, and there being agreement by the cognescenti (Devesh) that they were worthless I didn't need to do those worthless things.

Also on Fri, I worked with Devesh, who's been tapped as the contact-point in the Wahoo group for perf-related stuff. He had several interesting ideas about how to pursue the perf difficulties, all of which had come up in group discussion before, but which couldn't be followed up until now (because we didn't have a working Wahoo machine). I satisfied all his requirements, using PerfScore as "the" right tool for doing it. Devesh "gets it" how important PerfScore is to Netezza.

BTW, have you noticed that all the above were 5 days of work packed into only 3 days? I did this voluntarily, of course, as I always step up where above-and-beyond-the-call-of-duty is required. Nevertheless, that good deed didn't go unpunished, because Fritz shat upon me in public (Camb office) with lies, bullying/harassment and yelling, and surreptitiously (behind my back, refusing to talk to me face-to-face) causing me to be "fired" from the Wahoo project on Fri. This was an "illegal" adverse job action (in the IBM sense, perhaps even in the civil law sense), because it was a consummated false defamation of me (IBM policy calls it "harassment"), totally without due process.

The very act of Fritz's having gone to Dan behind my back, and falsely accusing me of being a liar and bully (as Dan told me he did)\*, already amounts to such an "illegal" act. Additionally compounding that defamation is the further public humiliation of unilateral removal from the most excellent high-profile position on Wahoo, to what seems (to me and others, just ask any

disinterested third-party observer) to be a highly symbolic deportation to Siberia. I felt unjustly accused and unjustly acted upon. Nevertheless, Dan reiterated his many-times repeated mantra that his ONLY [his word] interest is in helping the Wahoo project succeed (thereby excluding interests in such minor niceties as justice to me, if it came at the expense of Wahoo/Fritz). Upon telling Dan that I was interested in following-up this issue using established corporate processes, he suggested that I should go to legacy Netezza HR (Kelli-ann McCabe), and that I should also look up (on "IBM Blue Pages") the new IBM HR escalation processes. I have initiated those things. As time goes on, this activity will claim some of my hours to be devoted to it, which I will legitimately charge against hours I could have spent doing productive technical work.

Please observe: I am being straightforward, above-board and highly communicative about this whole adverse job action thing. This is in complete contrast to the way I myself have been treated (by Fritz, but also by Dan who has refused multiple requests by me for 3-way meetings with Fritz). I believe I am doing it "The IBM Way".

\* Footnote: That part of the meeting with Dan in his office Fri morn went like this: I reminded Dan that at our earlier meeting regarding a similar unaccountable action by Fritz against me (~ 3 weeks previously, where Fritz complained about my non-production of some graphs he had in fact never asked me to produce), I had stated that some things Fritz had said to Dan (as reported by Dan to me, recalling that Fritz had again gone behind my back to talk to Dan) were "provably known false when stated", but that I was hesitant to call them "lies" because to me that word had a connotation of "intent to deceive". But this time, after having reviewed all the facts and context available to me, I had come to the conclusion that I must indeed accuse now Fritz of the very word "lying", in addition to reiterating my previous charge of bullying. At that point Dan looked at me squarely and told me "I must tell you that Fritz has come to me and told me the same things about you".

## A.ee Week Of Jun 12-19

Spent the bulk of the week on The Unpleasantness. I'm charging that to company time, because the company (in the form of agents acting in their organizational capacity) is forcing me to do it.

Worked with Sujatha on "the transition", including machines. I lent the Skimmer to Devesh for most of the week, to do Wahoo-related testing. I hope that was OK, but most likely it isn't: I'll probably get docked for not hoarding it for myself and somehow thereby proving my worth, rather than letting Wahoo have all the glory.

Learned blktrace/blkparse/btt. The first 2 were introduced to me by Sujatha, but she never even looked at btt. It seems she simply grabbed the pre-compiled object files from a Snap server, rather than doing the research of obtaining the open-source code and compiling herself. Yet, btt seems to be the most useful/important of this suite. By far. Had to do all this on laptop, because of aforementioned loan of Skimmer to Devesh.

Put a bunch of stuff I've done on the Wiki (csvBeautify, nzMugshot, PMtest, PerfScore). I've been waiting to do this because Dan has been saying for many months he's looking into a reorganization of the Wiki to determine the best place to put stuff like this. But it's finally become obvious he has no intention of ever doing that.

Blockers: Being persecuted.

## A.ff Week Of Jun 19-26

Finished blktrace, published it on Wiki.

Talked to Sujatha about Emu.

Fri lost because no machine (colo down for maint), so can't actually look at Emu yet.

Fighting like hell to escape from abusive mgrs.

## A.gg Week Of Jun 26 - Jul 3

(Even though a weekly report was not required to be written this week, I'm doing it because I like having things on the record.)

Tried to begin Sujatha's EMU/FPGA project, but Garth's kit (which Sujatha had used) wouldn't boot. The symptom was weird: a certain executable complained about not being able to find a certain symbol in a certain shared library (identified by full pathname), even though the symbol was manifestly there (according to "strings", couldn't use "nm" because the library was stripped). Sujatha & Felix couldn't explain what was wrong. On a hunch I traded Felix for a different P50, and then things magically started working. Nobody could figure out what the problem was, but I lost a day because of it. My best guess is that there was a change of host RHELs between NPS revs that somehow broke which builds could run on which machines. Whatever.

Once I did then start looking at the EMU/FPGA project, it turned out there was no "there" there. This project had nothing to do with EMU/FPGA, it was only a trivial script monkey job, to capture a certain virtual table, `_vt_emulator`, that Garth had invented for his work on EMU/FPGA. But the scripts in use for that capturing were, uh, stupid (this is not "ad hominem", it's "ad codinem"). My recollection is that Sujatha had brought this project up at a group mtg, and somebody (or two) offered to help her. Be that as it may, however these scripts came into being were so bad that it raises serious questions about the usability of the data they had captured (because of the Heisenbug effect). So I wrote a proper program to do the job, generalizing it to all virtual tables in the process, and published it on the Wiki ("Capturing Complete Virtual Tables"). This should help numerous people, because I know that at least Devesh and Joseph had need for such a program. So that was the end of that project.

However, even though that was one of the best/fastest resolutions of an important problem I've seen for awhile, Dan had some sort of complaint about the way I did it. He said I should have asked him for direction about how to manage resources, rather than spend a couple of days fixing things once and for all. I didn't "get" his drift (I have my suspicions about what was going on, but I decline to state them here, see Metzger paragraph below).

I got surprised by having last week's blktrace project as Sujatha had described it to me (another script monkey project, with Dan as the customer instead of Garth) changed out from under me. The new thing is to actually do the analysis work myself (which is much more interesting, just new/different from what I'd been told the project was by Sujatha). That means I'll now have to drop/delay the oprofile project indefinitely.

Had a talk with John Metzger, about "cooling it" regarding "tone" (as opposed to "content", no one has yet seriously challenged that) of my emails (even though the precise emails he was referring to were supposedly under the protected aegis of HR process). He suggested limit-



ing myself to objective facts, and suppressing subjective opinions. I respect John, his handling of this has to date been a refreshing "steady hand on the tiller" approach, and I thought this was a wise suggestion, so I agreed I would try hard to do that.

Started the IBM Concerns and Appeals process, w.r.t. The Unpleasantness. This will take a non-trivial amount of my time going forward, because it appears to be much more rigorous, especially involving the write-up of a full-fledged "complaint" (in the style of a regular complaint at civil court). Therefore I will properly charge the time to the company, since this is an official company process, and because it's been the company itself all along that has proactively caused the whole thing to happen, and refused to deal with it in a straightforward/correct manner (as advertised in all its propaganda, such as employee handbooks, about how good/humane IBM is in treating its employees). And yes, I believe this is indeed an objective statement, as I will continue trying to prove until I have no avenues for redress remaining. I may lose, but at least I've have done it the right way.

## A.hh Week Of Jul 3-10

(2-day week: Mon = 4th of July; Thu&Fri = medical leave.)

Only real work this week was on blktracing.

Talked with Garth about his ideas for blktracing. His first suggestion was to look at read-only, so I started doing that, but then at a later mtg with Garth & DanF that was reconsidered to include writes from the beginning as well. So what I ended up doing was to start looking at all 3 combos from the beginning (r+w, r-only, w-only). It was also agreed to start looking at multi-stream muqry from the beginning (skipping an earlier thought to look at single-stream Atomics+TPC-DS).

Design & impl of nzBlkParseFilter.py, a tool necessary for analysis (of the kind desired, see above) of blktrace (i.e., this wasn't already covered by the regular blktrace/blkparse/btt suite).

Ran Muqry@15000B, during which captured ~14.5 min blktrace, and analyzed it (to the extent we currently know how to do that, i.e., what to look for). Produced some graphs (including very-wide-landscape 11"x25'), thus getting some idea of what multi-stream looks like when it hits the disk. Published it all on wiki. Pretty interesting (e.g., seemingly unexplored previously), though still unclear how useful this will really be in the long run for improving NPS perf.

2 hrs going back and forth with DanF in email about how to write the kind of "plan" he wants. I still don't get that.

Nearly 2 hr in mtg with DanF, trying to clear the air, but still a bit murky. But we both agreed (didn't we?) to focus on "the future, going forward", though as for the past I've made it clear I still don't have closure on that and am pursuing it (IBM "Appeals and Concderns" process).

Next 3 weeks: Medical leave for 2 weeks, then vacation for 1 week, returning on Aug 1. Hence, am including the relevant Weekly Reports here:

## **A.ii Week Of Jul 10-17**

Out on medical leave (short-term disability).

## **A.jj Week Of Jul 17-24**

Out on medical leave (short-term disability).

## **A.kk Week Of Jul 24-31**

Out on vacation.

## **A.II Week Of Jul 31 - Aug 7**

Worked on nothing but "blktrace stuff" the whole week. This is in quotes, because the project is ill-defined currently (nothing wrong with that, just makes it harder to talk about), ranging from a low-end of "blktrace tools and training", to high-end of "figuring out how NPS uses its storage devices. At this point, I've more-or-less finished the low-end, and made some beginning inroads into the high-end. All this is getting documented on the "middle-end" wiki.

## **A.mm Week Of Aug 7-14**

Worked the whole week at home, because didn't feel well enough to travel to Marl (upset stomach, can't eat, can't sleep, unsafe to drive, but mind is fine so able to do Netezza work).

But Dan instructed me not to work on blktrace, and I have nothing else to do either (due to previous instruction not to work on oprofile, and finished with nzVtCapture.sh, and Sujatha's finishing up FPGA runtime restricts), so I did no work at all. The first 4 of these days (Mon-Thu) were "free", in the sense that I was reading/willing/able to work (at home), but Dan told me not to, so I didn't. But I decided to take Fri as a sick day anyway.

Also, started the process of applying for STD and workplace accommodation. Parts of that, at least, should be in place by next week.

## **B Email Chain: Database Transport (March 15-16)**

■ From: Devesh Agrawal  
To: Walter Tuvell  
Date: 03/15/2011 9:52 PM  
Subject: Safe to do updates/deletes?

Hi Walt -- I think I have managed to make deletes and updates do nothing more than a scan. So they should be non destructive for Wahoo now. Try it and let me

know if thats indeed the case. Thanks. { So you can just load your simple Bar loading scripts instead of worrying about how to get the data into the wahoobox }.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 03/16/2011 9:21 AM  
Subject: Re: Safe to do updates/deletes?

Will try, thx!!

■ From: Devesh Agrawal  
To: Jerrold Richard Title, Walter Tuvell  
Date: 03/16/2011 10:25 AM  
Subject: Wahoobox ...

Hi Folks,

Please keep in mind that the wahoobox won't work with a pre-merge kit anymore. Since the initTopology files have been changed to the new version. So Walt -- you might want to checkout and build your own kit of dev-wahoo instead of Rich's old kit. After that it should come up with 4 slices.

■ From: Walter Tuvell  
To: Fritz Knabe  
Bcc: Daniel Feldman  
Date: 03/16/2011 10:45 AM  
Subject: Testing status

Fritz, as we've discussed, there's no current build for WahooBox that can create a TPS-DB database, because of the DELETE crash problem left by James (i.e., as his instigation to Devesh to make some change that causes the problem).

Last night Devesh checked-in a fix that purportedly maps DELETE and INSERT to no-ops (really, table-scans).

But at the same time, the merge is still in flux. So building Wahoo at this point won't work, until that gets straightened out (Rich, Jeff & Co.).

I'm also proceeding with the DB-copying scheme. See below for current progress. Even so, there's no guarantee this will work in any sense, such as will it be importable to WahooBox (we know we can't just do the regular B&R, and first by-hand efforts are always error prone)? And there would certainly be collisions between me and others (esp. Rich) when it comes to trying to use WahooBox the next few days.

Does this really seem like a good use of time? I've gotten involved in several productive projects during Wahoo's absence from usability, but this doesn't seem like one of them.

```
-r--r--r-- 1 nz nz 29G Mar 15 22:05 CATALOG_SALES.csv
-r--r--r-- 1 nz nz 142M Mar 15 22:06 CUSTOMER_ADDRESS.csv
-r--r--r-- 1 nz nz 385M Mar 15 22:05 CUSTOMER.csv
-r--r--r-- 1 nz nz 100M Mar 15 22:06 CUSTOMER_DEMOGRAPHICS.csv
-r--r--r-- 1 nz nz 9.9M Mar 15 22:06 DATE_DIM.csv
-r--r--r-- 1 nz nz 100 Mar 15 22:06 DBGEN_VERSION.csv
-r--r--r-- 1 nz nz 194K Mar 15 22:06 HOUSEHOLD_DEMOGRAPHICS.csv
-r--r--r-- 1 nz nz 308 Mar 15 22:06 INCOME_BAND.csv
-r--r--r-- 1 nz nz 7.7G Mar 15 22:20 INVENTORY.csv
-r--r--r-- 1 nz nz 100M Mar 15 22:20 ITEM.csv
-r--r--r-- 1 nz nz 173K Mar 15 22:20 PROMOTION.csv
-r--r--r-- 1 nz nz 6.5K Mar 15 22:20 REASON.csv
-r--r--r-- 1 nz nz 2.1K Mar 15 22:20 SHIP_MODE.csv
-r--r--r-- 1 nz nz 118K Mar 15 22:20 STORE.csv
-r--r--r-- 1 nz nz 3.3G Mar 15 22:29 STORE_RETURNS.csv
-r--r--r-- 1 nz nz 38G Mar 16 00:20 STORE_SALES.csv
-r--r--r-- 1 nz nz 7.6M Mar 16 00:20 TIME_DIM.csv
-r--r--r-- 1 nz nz 2.2K Mar 16 00:20 WAREHOUSE.csv
-r--r--r-- 1 nz nz 55K Mar 16 00:20 WEB_PAGE.csv
-r--r--r-- 1 nz nz 996M Mar 16 00:22 WEB_RETURNS.csv
-r--r--r-- 1 nz nz 15G Mar 16 01:01 WEB_SALES.csv
-r--r--r-- 1 nz nz 26K Mar 16 01:01 WEB_SITE.csv
-rw-rw-r-- 1 nz nz 3.5K Mar 16 09:24 CALL_CENTER.csv.tgz
-rw-rw-r-- 1 nz nz 462K Mar 16 09:24 CATALOG_PAGE.csv.tgz
-rw-rw-r-- 1 nz nz 929M Mar 16 09:31 CATALOG_RETURNS.csv.tgz
-rw-rw-r-- 1 nz nz 8.6G Mar 16 10:34 CATALOG_SALES.csv.tgz
```

■ From: Jerrold Richard Title  
To: Devesh Agrawal  
Cc: Walter Tuvell  
Date: 03/16/2011 10:47 AM  
Subject: Re: Wahoobox ...

Walt doesn't have to do his own dev-wahoo build, unless he wants to.  
He can use my new kit after I test it out.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal  
Date: 03/16/2011 10:56 AM  
Subject: Re: Wahoobox ...

That would be great, because I haven't done a build since my first month here, and I've have to review my notes about how to do it. I'd also have to ask how things have changed since then, and how to configure core affinity and number of slices,

things of that nature. You guys are always helpful of course, and it's not rocket science, but the laziest path is always the best!

BTW, Rich, I need the fixes Devesh checked in yesterday (otherwise TPC-DS database creation crashed), would that be in your kit?

■ From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Devesh Agrawal  
Date: 03/16/2011 11:18 AM  
Subject: Re: Wahooobox ...

Yes, Devesh's fixes are checked into dev-wahoo and would be in my new kit.

Restart (nzstop/nzstart) will be broken in the new kit. As I said in my earlier email, I expect that to be fixed in a day or two. If you need restart now, you're SOL. :-)

Tomorrow when you're in the office, Devesh or me can sit down with you and go through the steps to create and build a workspace. It's not too hard. :-)

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal  
Date: 03/16/2011 12:22 PM  
Subject: Re: Wahooobox ...

What does "start/stop is broken" mean? Would the DB have to be reloaded from scratch, for example?

If "all goes well", the whole TPC-DS test could be done in a single shot, without restart in-between.

In any case, it would be my preference to wait until your stuff is working. The real question is what Fritz's preference is, and I feel he's sitting on me at the moment (not making a good choice there, IMHO).

■ From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Devesh Agrawal  
Date: 03/16/2011 03:19 PM  
Subject: Re: Wahooobox ...

Hi Walt,

That's exactly what it means - you have to do your test in a single shot, or reload

your data.

I expect to have stop/start functioning again in a day or so. Dan is helping with this.

I verified that Devesh is correct (of course) - the old kit no longer runs. I assume that's because his start-up scripts on wahoobox are in sync with the new kit, not the old kit.

I have left wahoobox running on my "turbo" kit and I am not using it for now. You can have it, if you can live within the no-restart constraint.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal  
Date: 03/16/2011 04:25 PM  
Subject: Re: Wahoobox ...

Sure thing, I'll take it, thx!

■ From: Walter Tuvell  
To: Cambridge  
Date: 03/17/2011 12:44 PM  
Subject: PerfBar test results

As mentioned at standup, WahooProto got through a partial run of TPD-DS (100GB) last night. The results are below, compared to Skimmer.

As always, numbers mean seconds per test, and negative numbers mean the test failed for whatever reason.

The numbers aren't good. Are we sure all fixes are still in the build (i.e., no merge accidents)?

QUERYID	SKIMMER	WAHOOPROTO
query001	10.33	85.47
query001v	8.88	139.61
query002	81.64	894.62
query002v	38.80	511.71
query003	16.14	634.67
query006	66.84	724.73
query007	17.28	658.31
query008	17.80	616.80
query011	124.09	1059.20
query011v	79.48	1701.83
query012	9.11	178.97
query013	20.38	717.23
query014b	-65.80	-1947.63
query014bv	-0.17	-0.17



query015		26.80	412.74
query016		23.98	407.96
query017		-0.19	-0.16
query018		17.43	255.58
query019		19.13	657.06

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Cambridge  
Date: 03/17/2011 01:57 PM  
Subject: Re: PerfBar test results

Those certainly are disappointing numbers. When you have the data with what the CPU utilization was versus I/O traffic over the course of the test, that will be helpful. Right now this looks like an "Oops!" problem with some component that can hopefully be pinpointed quickly.

## C Email Chain: Excel Graphics (May 19)

■ From: Walter Tuvell  
To: Daniel Feldman, Fritz Knabe  
Date: 05/19/2011 7:37 AM  
Subject: Out today

The deck work has been postponed, but instead I have 2 Dr. visits today, and too I'm feeling sick-of-something, so I won't be in.

Fritz, I guess this means I'll be missing the iteration planning (which you scheduled last night at 8 pm and I just now saw).

← 2 additional emails inadvertently omitted; New Complaint, Add. III, bullet-item spanning pp. 7-8

■ From: Fritz Knabe  
To: Walter Tuvell  
Date: 05/19/2011 11:37 AM  
Subject: Re: Out today

For the iteration planning, the big push for performance will be building up the relative perf stats for Wahoo vs. TwinFin (Skimmer) on TPC-DS.

Sorry about the surprise on scheduling! I've been mentioning it at stand up a few times this week, but indeed hadn't sent out the formal invites.

■ From: Walter Tuvell  
To: Daniel Feldman

Date: 05/19/2011 11:49 AM  
Subject: Fw: Out today

More mind-reading. I'd already scheduled today and tomorrow off \*LAST\* week, his mention of iteration planning this week was twice (not a "few"), and the word-ing was "sometime this week", which in the past has been Thu \*OR\* Fri. And it was probably the only interesting thing that was said at stand-ups, which are a complete waste of time, as is the planning meeting.

Beside which, if he wants "relative stats", how is that supposed to happen if there doesn't \*EXIST\* a \*TOOL\* to do it, or even a \*COHERENT CONCEPT\* of what the idea even is??? [Now there finally is, PerfScore.]

Somebody's brain and/or ego isn't working right, it's not you or me, and I'm going to need managerial help on it.

■ From: Daniel Feldman  
To: Walter Tuvell  
Date: 05/19/2011 12:11 PM  
Subject: Re: Fw: Out today

Yes, I agree that the PerfScore work provides a defensible mechanism for comparing between the two architectures and is valuable and necessary work; it is good that it is ready to go.

I'm not too concerned about you missing his iteration planning meeting; you gave him plenty of notice and he was imprecise about scheduling it.

I'll think about how we get this onto a sustainable, professional footing. I'm certainly open to all serious suggestions.

■ From: Walter Tuvell  
To: Daniel Feldman  
Date: 05/19/2011 6:11 PM  
Subject: Re: Fw: Out today

About the only thing I can think of is the "obvious": We need to have a/another 3-way talk. I/(we?) can't let that slip.

<flame>  
A bully can't be ignored, else they'll keep on bullying.  
</flame>

## D Email Chain: Wahoo Not The Bottleneck (May 31)

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/31/2011 04:29 PM  
Subject: Fixing nmon

Nigel Griffeth (IBM UK, owner of nmon) got back to me over the weekend. He said my request was new, gave me some hints about how to do it, and said both that he thought the feature didn't belong in nmon so I should do it myself, and that he'd think about putting it in(!). I didn't have the heart to tell him that I'd already done the deed, better than the way he suggested (nanosleep() vs. usleep()), the latter is obsolete because it doesn't tell you how much your sleep got awakened by a wayward signal so you can sleep a little longer if necessary, and yes this is in the code). He put the conversation on his forum, and you can see the conversation here: <https://www.ibm.com/developerworks/forums/thread.jspa?threadID=372385&tstart=0>.

Even at the time, I recognized the problem of skew: if you sleep 0.1 sec each time, you'll drift, thereby sometimes getting only 9 updates in some seconds (maybe this could have been messing up some of our earlier data, because Devesh's perfMeasure didn't deal with that?). But I didn't know how to fix it correctly, specifically I recognized that only doing part (i) of the following sentence isn't sufficient. I have now today fixed the problem, by: (i) noticing how much too long/short each sleep was, and dynamically micro-adjusting accordingly; (ii) basing the overall structure on the number of update intervals, as opposed to timestamp readings per se. This "guarantees" (to the accuracy we care about) that our nmon measurements will be "good" (or at least "as good as possible not using real-time facilities", which I think is "good enough").

All this is hack hard-coded into nmon, specifically, only 0.1 sec is supported. (Certainly I could have done the general job, but that requires polishing I didn't want to invest.) So my version has a new name: nzmon.

Devesh tells me Violin supports its vcount stats via /proc, so I'll put that into nzmon too,

I have this running on host (easy, as suspected), will ask Larry/Felix how to get it running on SPU (DanD didn't have a ready answer, apart from putting the code into Accurev, it's unclear whether that's desirable).

■ From: Daniel Feldman  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 05/31/2011 4:38 PM

Subject: Re: Fixing nmon

Nice, thanks!

Larry did some work to run nmon on blades (perhaps using nzipush with the "secret" switch that lets you push anything you want, but I'm not sure). The "big" problem is what to do with the data from n blades of m CPUs each. I think this is a place where R could help but we might want to spend some time thinking/talking about what the output would/should look like.

Also, as long as we're in nzmon-land, Joe has invested in some /proc based per thread data collection. Might be something to think about including at some point; seems off-topic for right now, though.

## E Email Ensemble: Continuing Untestability Of Wahoo (May 23 - June 10)

►*The emails in this Appendix provide copious documentary evidence for the ongoing unavailability of Wahoo during the critical period May 23 — June 10. But there were even additional issues that were never documented, because they were merely communicated by word-of-mouth.*◄

■ From: Walter Tuvell  
To: Daniel Feldman  
Date: 05/23/2011 09:45 AM  
Subject: Tasks for this iteration

For this iteration Fritz wants a comparative study of Wahoo and Skimmer on TPC-DS (100GB first, then 1000GB if there's time).

To complete this assignment, I plan to first get a set of numbers from Skimmer, then from Wahoo, then compare them using the perfScore.py tool.

To get the numbers from Skimmer involves "porting" TPC-DS to a non-ballooned DB (necessary, because Wahoo doesn't yet support ballooning). My recollection (from ~ 6 weeks ago), is this isn't a major problem, but I could be wrong. The output of this task is the PerfBar "summary" containing the required numbers. This should be finished this week.

To get the numbers from Wahoo involves running that (now-modified) TPC-DS on Wahoo. If past experience is a guide, some additional problem will be hit, esp. bugs in Wahoo (for example, it may fail on some tests that Skimmer didn't, despite the non-ballooned nature of the DB). The output is again a PerfBar summary file. This should be finished next week.

Running perfScore.py against those two summary files should take only a very small amount of time (transferring the summary files to my laptop, running the tool,

packaging and communicating the results). Should also be finished next week.

The 1000GB analog of all the above will be done the third week (which ends on Thu Jun 16 for the purposes of this iteration).

This plan assumes there will be various problems and diversions along the way. If those don't materialize, the job will be done earlier, but I don't anticipate any delays will extend the plan beyond the stated times.

■ From: Devesh Agrawal  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 05/24/2011 11:49 PM  
Subject: Lates on wahoobox's generate stats problem reported today

I was able to load a 1 and 10G Tpcds and do genstats on it (by hand without using the 401\_create script). I am sure the same would work with even a 100G dataset -- I didn't try it since I ran out of hard disk space on my workstation when generating the dbgen dat files.

But when I did a 10G tpcds load with the 401 script it went through the stage of generate stats but crashed the SPU when it tried to do a "delete" operation from the fact tables to remove the NULL columns. Earlier, delete was a harmless operation but I think Rich changed that so that we now do some half baked delete which goes and crashes the SPU. :).

So my advise is -- lets not do the 401 thingy, I think all it gives us is the ability to do is parallel load streams. Or one can modify the 401 script to not do anything except nzload and genstats.

Wahoo1 currently points to my kit but please change it to yours when you test it.

Thanks.

■ From: Walter Tuvell  
To: netezza-cambridge  
Date: 05/25/2011 01:54 PM  
Subject: Update on progress

Here's a little update on the progress of TPC-DS testing. See the included screenshot. The titles on windows show what machines they belong to, and the "tower" you may not recognize is gkrellm (on the hosts, not SPUs).

On Skimmer (S-8), life is good. It has now run 2 runs at the 1000GB, and is in the middle of a third (at query048). No problem.

Wahoo1/2 are having problems. On Wahoo1, I disabled the GENERATE STATISTICS we thought was causing problems (that's the yellow line highlighted in the emacs window). As you can see, Wahoo1 is working on query011. As for Wahoo2, it is stuck on GEN STATS, just like it was yesterday.

But that's not the real story. The real story is that Wahoo1 is stuck too. What we're seeing is that sporadic progress is being made by Wahoo1, but some of the queries are super-slow, such it looks "as if" the machine is stuck. In other words, GEN STATS may not really be a problem at all, it may simply be one of the super-slow queries.

I'm going to let this run for awhile longer, to get as much information as I can from this test, then try some new things. After all, since this is the first test on DwarfHoo, the problem may be related to that, so a first thing to do is try backing-off from the simultaneous tests, to a single test, then all the way back to physical WahooBox if necessary. (This with Steve's help, of course.)

After that, probably re-running with some more monitoring/instrumentation enabled seems like a good idea, to see if can get a better idea what the machines think they're doing at the time of the super-slow queries.

On the positive side, notice how nicely sync'd all the clocks are!

Be talking to you.

►*Screenshot omitted here; irrelevant to this Complaint.*◄

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Steve Lubars  
Date: 05/25/2011 04:00 PM  
Subject: Re: Update on progress

Thanks for the update. Since we're new to the VM setup on the prototype, I would be inclined to do as fast a test as possible on the physical configuration to quickly determine whether we're dealing with VM issues or with general stuff. I assume you're testing on the 1 GB data set, as we discussed yesterday, so that will hopefully load and run pretty fast if everything works. Steve, how painful is it to make the switch from one configuration to the other? We can consider a different path if it will take too many hours to do that switch.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Steve Lubars  
Date: 05/25/2011 04:14 PM  
Subject: Re: Update on progress



We already did the transition 10 min ago, it's booting now.

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Steve Lubars  
Date: 05/25/2011 04:22 PM  
Subject: Re: Update on progress

Great! I'm not sure which result I would prefer to see, but positive or negative, this will tell us something.

■ From: Devesh Agrawal  
To: Jeffrey Keller  
Cc: Daniel Dietterich, Walter Tuvell  
Date: 05/25/2011 08:11 PM  
Subject: Wahoo12 crashed in addPrecommittedTxid

addPrecommittedTxid call that you added probably needs to do some locking before updating the STL map since it is shared by the whole DSDesc. There was probably a corruption caused in there. Here's the stack trace:

Walt, you might want to try loading serially instead of the default num-cpu's parallel streams that 401 stream does. Check out the serial loading option that will make your load 4 times slower, but at least work hopefully. Sorry for this.

```
at /home/workspaces/wtuvell/wahoo12/main/obj/crosscc/base/bin/./lib/gcc/i586-wrs-linux-gnu/4.3.2/../../../../i586-wrs-linux-gnu/include/c++/4.3.2/bits/stl_tree.h:1153
```

```
#7 std::set<TXid, std::less<TXid>, rc::StlAllocAdaptor<TXid, NzShmAlloc<TXid, 192>>>::insert ( _x=<value optimized out>, this=<value optimized out>) at /home/workspaces/wtuvell/wahoo12/main/obj/crosscc/base/bin/./lib/gcc/i586-wrs-linux-gnu/4.3.2/../../../../i586-wrs-linux-gnu/include/c++/4.3.2/bits/stl_set.h:381
```

```
#8 pm::DsDesc::addPrecommittedTxid (this=0x704f8224, txid={value = 8582424828}) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/pm/pmDsDesc.cpp:163
```

```
#9 0x083e2e33 in pm::Tx::txOperation (this=0x73bcd7a0, txop=pm::TXOP_PRECOMMIT, dsdesc=0x704f8224) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/pm/pmtx.cpp:288
```

```
#10 0x083e3c14 in pm::TxIndex::txOperation (this=0x705626c8, txop=pm::TXOP_PRECOMMIT, txid=1190, dsdesc=0x704f8224) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/pm/pmtx.cpp:447
```

#11 0x083e58a0 in pm::DsMgr::txOperation (this=0x705626c0, txop=pm::TXOP\_PRECOMMIT, txid=1190) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/pm/pmDsMgr.cpp:45

#12 0x083e0d05 in pm::ADSHand\_::precommit (this=0x8fb07ddb, txid=1190) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/pm/pmAdsHand.cpp:26

#13 0x0819cc93 in CSpuTXMgr::finishPrecommitTx (this=0x700ce320, xId=@0x75feddc4, dataslice=4, from=0x70565904) at /home/workspaces/wtuvell/wahoo12/main/src/nde/acid/SpuTXMgr.cpp:802

#14 0x0814064a in handleInJobTask (arg=0x723eebb0) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spumain.cpp:3225

#15 0x081a5cc0 in \_invoke\_entry (arg=<value optimized out>, entry=<value optimized out>) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:737

#16 \_childproc (jobtask=0x750d4184, priority=10, time\_slice=10, pipe=@0xf8d6db0) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:812

#17 0x081a5e9f in \_\_create\_task (jobtask=0x750d4184, priority=10, time\_slice=10, pipe=@0xf8d6db0) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:851

#18 0x081a5fe9 in \_create\_tasks (req=0x70563d74, priority=10, time\_slice=10) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:916

---Type <return> to continue, or q <return> to quit---

#19 0x081a62e1 in \_create\_job\_tasks (req=0x70563d74) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1294

#20 0x081a73c3 in nzprocmgr\_main () at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1753

#21 0x08141dce in NzSpuMain () at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spumain.cpp:1054

#22 0x08142436 in main (argc=Cannot access memory at address 0x0 ) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spumain.cpp:1064

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/26/2011 02:32 PM  
Subject: Status

I have to leave now (prom night, daughter needs car), will finish up from home.

The 100G on Wahoo is still(!) loading.

Nmon looks ideal: simple/portable C program, reads /proc, produces curses GUI or CSV files. Got it running on my laptop in no time. Perfect.

Except for one small(-ish?) problem: It only supports full-second update granularity (hence correspondingly only full-second timestamps). Being open-source, I've looked at the code, and there's some trickishness that make sub-second granularity slightly non-trivial, but not rocket science. All-in-all, this looks like the best path to start down. Simultaneously I'll contact the developer, and see what he has to say (it's currently active, last updated in Jan). He works for IBM and does nmon on-the-side, but maybe the filial relationship will have some sway?...

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/26/2011 08:38 PM  
Subject: Status

STILL building 100G DB at this hour, but I keep early hours, so I've set a cron job for running TPC-DS, hope it doesn't have a typo (and that the DB has built by 1am).

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Daniel Feldman  
Date: 05/26/2011 09:16 PM  
Subject: Re: Status

Jeff has checked in the fix for the bug. If it doesn't finish before we're at work, we may want to kill it, allowing us to do any of the following:

. Use the new build to run a 1GB sanity-check test and then a 100GB test

. Or if networking instrumentation is done, proceed with using the prototype machine to test out the host CPU utilization options, and then get to the 1GB and 100GB tests.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Daniel Feldman  
Date: 05/27/2011 09:10 AM  
Subject: Re: Status

It's STILL loading at this time (really loading, not hung). So our "8x" estimate for 1-thread vs. 8-thread was off by a factor of > 4.

By analyzing the logfile, I estimate we're ~90% loaded, should be done by ~noon or shortly thereafter. The principle of "preferring the devil you know to one you don't" would argue that we continue along this path, because it gives the greatest likelihood of getting some results today (e.g., how do we know Jeff's bugfix works, or that it's the only bug in parallel loading?).

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/27/2011 09:32 AM  
Subject: Fw: Status

Just got confirmation that plunging forward it the only tenable plan right now: Devesh royally broke dev-wahoo last night ...

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/27/2011 01:04 PM  
Subject: Latest

See screenshot below. WahooBox (which got amnesia, and now thinks its name is "localhost") is steadily churning away (CPU2), throwing 0.5-1.0 MB/sec of DB at the Fender. Currently there are > 221M rows in the store\_sales table, on the way towards 287M (shown on Skimmer S1-8's window on the right). After store\_sales is full, all that's left is web\_sales, which has 72M rows.

So progress is happening, but I wouldn't call it "steady". It actually seems to be sub-linear, slowing down over time. Good thing we've got a long weekend coming up ...

►*Screenshot omitted here; irrelevant to this Complaint.*◄

■ From: Daniel Feldman  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 05/27/2011 02:22 PM  
Subject: Re: Latest

The slowing down phenomenon seems problematic, doesn't it? Any ideas on how we might diagnose it while we're well into a case like this?

■ From: Walter Tuvell  
To: Daniel Feldman  
Cc: Fritz Knabe  
Date: 05/27/2011 02:31 PM  
Subject: Re: Latest

Agree it seems problematic, and I mentioned at stand-up there could be some bug beyond the parallel-load bug (sorry, I could have mentioned that to you). However, I'm not so sure I want to start down that path without first seeing what the full-parallel-load looks like. For, I haven't actually been monitoring per-minute loading, and I could be imagining things -- the load has taken much longer than I expected, and maybe I'm just getting antsy. At this moment: 253M rows out of 287M, with the 75M-row web\_sales table then to go.

Another consideration is prioritization: I have got nmon sub-second samples working, but not have to get sub-second timestamps working. (Wrote to Nigel Griffiths, nmon Daddy, he hasn't gotten back to me yet, if ever.)

Of course, that's just my priority, I could be over-ruled ...

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/31/2011 06:16 AM  
Subject: Where we stand

WahooBox is still broken.

Recall: Last week the best we were able to do was create a 1GB and run PerfBar on it, which is OK for basic qualification purposes, but too small to take seriously as a comparative, or report at a mtg.

We've also been able to create a 100GB, but when trying to run PerfBar we've been hanging. The conjecture (assuredly partially correct) by Devesh was that DB corruption was happening at DB-creation time, and he suggested tweaking several settings to avoid that, one of them being "--parallel=1" to avoid Jeff's bug, plus a few others. I've adopted all the suggestions, and even tried combinations thereof, but to no avail. We're still hanging at the PerfBar stage. Hangs have occurred at places like query011, query012, query014b, query019 (out of ~100).

Most recently (yesterday) I tried cutting the baby in half, and try a 10GB, but that hung during the DB-creation stage. Interestingly, I created 3 different 100GBs with "--parallel=1", and it took 36 hrs, 8 hrs, 4 hrs respectively, even though I was starting from totally fresh reinit each time (there were some slightly different settings as mentioned above, but certainly not different enough to account for these time differences). Go figure.

This seems fruitless. We need to fix Wahoo before we can get measurable numbers.

■ From: Daniel Feldman  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 05/31/2011 06:57 AM  
Subject: Re: Where we stand

Thanks for the update and the work over the weekend. This is interesting in its own way although it does slow us down on getting to the kind of comparison Fritz is looking for. I don't think the changes in load times for the same operation can be easily dismissed. Can you create a table that shows, for each of the three 100GB load attempts, the kit you were using, the settings of any switches/parameters/configuration variables that were used during the load, the date and time the load started and the elapsed time of the load? I think there is a slim chance that someone might be able to diagnose or usefully begin to diagnose the cause(s) of the different load times with that table as a starting point.

■ From: Walter Tuvell  
To: Daniel Feldman  
Cc: Fritz Knabe  
Date: 05/31/2011 09:46 AM  
Subject: Re: Where we stand

The answer, insofar as I can give it, is below.

Right now, what I've done is re-gen the 1GB, and running PerfBar on successively, trying to see if it will hang even on a 1GB at some point. Currently in middle of 1st run, no hang yet. Of course, this won't yield any useful numbers for us, only potentially debugging info.

```
401_create_tpcds_db -s 100 --noballoon -parallel=1
```

Three DB creations of TPC-DS@100 DB, all using the above options (plus others, see below), took times of 36 hrs, 8 hrs, 4 hrs. In all cases, the 401 script succeeded, but runBar.py --test=tpcds failed (hung).

All creations were from a fresh install of NPS (restart\_wahoo.sh -S 1 -D 8 -R). The only differences amongst these 3 creations were different settings of the following:

```
genAllDbStats ON/OFF -- for us, this affects runBar.py, not the 401 script
updDsid ON/OFF
redistrib_web.sql ON/OFF
--skipstaging ON/OFF
```

But I didn't keep records (and now I forget) which settings belonged to which creations, except that: one of them used all settings OFF, and one of them used all settings ON.

In the above:



- 1) "genAllDbStats" is a function in lib/nz\_functions.py. To turn it off, insert a return statement at the beginning.
- 2) "updDsid" is a function in the 401 script. To turn it off, insert a return statement at the beginning.
- 3) "redistrib\_web.sql" occurs in 1 place in the 401 script. To turn it off, comment-out that line and its 2 surrounding lines.
- 4) "--skipstaging" is an option to the 401 script.

The wahoo kit used was the one from a May 23 update of dev-wahoo -DWAHOO.

■ From: Walter Tuvell  
To: Fritz Knabe, Daniel Feldman  
Date: 05/31/2011 11:00 AM  
Subject: Hung

As a follow-up: On the 1G db, the test is indeed hung now, on query077. And this is with avoidance of all the bad things Devesh suggested should be avoided.

■ From: Fritz Knabe  
To: Devesh Agrawal, Walter Tuvell, Daniel Dietterich, Jerrold Richard Title  
Cc: Jeffrey Keller, Daniel Feldman  
Date: 05/31/2011 01:26 PM  
Subject: Summary of bug discussion

Here are the actions:

- . Devesh will examine the currently hung wahoobox for core files, etc.
- . Rich will prepare a Wahoo stream rolled back in time before GC, async writing, and delete that Walt can use for another 1 GB TPC-DS test.
- . Walt will test this stream using a debug build. The hope is that if it does hang, the debug build will help with diagnosis. If it passes, we won't yet know for sure that the problem is not present in this stream, as it is nondeterministic and has passed 1 GB before. Subsequent actions will be to run 1 GB in turbo, then 100 GB in turbo.
- . Wahoobox is currently configured in physical mode. At some point here we will want to reconfigure to VM mode and assess whether there is any performance difference, and whether the two VMs affect each other's performance or are relatively isolated.

- From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Daniel Dietterich, Devesh Agrawal, Jeffrey Keller, Fritz Knabe  
Date: 05/31/2011 01:44 PM  
Subject: Re: Summary of bug discussion

I created a stream named usr-rttitle-wahoo-timebasis with a timebasis of Apr 27 6pm. I created a workspace under that stream at /net/rtitlerhl/home/workspaces/rtittitle-wahoo-timebasis/main/.... I am in the process of doing debug and turbo builds in that workspace (which will appear in the debug and turbo subdirectories when completed), and then I'll do some testing on the builds. You can use my workspace or create your own.

- From: Walter Tuvell  
To: Jerrold Richard Title  
Date: 05/31/2011 01:47 PM  
Subject: Re: Summary of bug discussion

I might as well wait and use yours when certified sane. The fewer moving parts, the less chance for human error.

- From: Daniel Feldman  
To: Walter Tuvell  
Date: 05/31/2011 02:57 PM  
Subject: Re: Where we stand

I realize that we didn't start out to measure load time and so the loading process hasn't been controlled with the same rigor as a query test would have been, but it would have been helpful to know what differed for each of the tests. Can you identify any change(s) between the 36 hour and the 8 hour test? (I realize that this may not be possible, just looking around for a clue that might lead us to an early diagnosis of the change from 36 to 8 to 4 hours.)

- From: Daniel Dietterich  
To: Walter Tuvell  
Date: 05/31/2011 03:17 PM  
Subject: Alter group hoping it works around the hang

alter group batch\_grp resource minimum 100 resource maximum 100;

- From: Jerrold Richard Title  
To: Walter Tuvell  
Date: 05/31/2011 04:26 PM  
Cc: Daniel Dietterich, Devesh Agrawal, Jeffrey Keller, Fritz Knabe

Subject: Sherman and Peabody's wayback machine

Well, leaving aside the question of whether you or I are certifiably sane, I did some sanity testing of these builds that put you back in time to April 27:

/net/rtitlerhl/home/workspaces/rtitle/wahoo-timebasis/main/debug

/net/rtitlerhl/home/workspaces/rtitle/wahoo-timebasis/main/turbo

I tested on VSPU, and on vm-dw88.

Feel free to give either one a whirl on wahoobox to see if they exhibit the problem running TPCDS that you saw with more recent builds.

Let me know if you want any assistance.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Daniel Dietterich, Devesh Agrawal, Jeffrey Keller, Fritz Knabe  
Bcc: Daniel Feldman  
Subject: Re: Sherman and Peabody's wayback machine

Thanks Rich, there's a fairly good chance I'll end up using this, but at the moment I'm working on DanD's conjectured fix, which is to adjust group resource allocations. To that end, I'm now build the 100G DB, and should be able to test it tonight (unless Wahoo decides to throw another 36 hr build, but I'm monitoring it and it seems to be having one of its better builds).

■ From: Devesh Agrawal  
To: Walter Tuvell  
Cc: Daniel Dietterich ►*Forwarded to Dan Feldman, 06/31/2011 10:04 AM.*◄  
Date: 05/31/2011 09:56 PM  
Subject: Wahoobox is down

I am not sure why it cored. The stack trace suggests this may be nothing to do with us (I hope !). I am wondering if there is a memory limit somewhere (artificial for sure) that is causing this to fail. I am not so familiar with this hashing/DHJ code to comment further.

(Ps: I am wondering if the 1GB run succeeded through ?)

from /nz/kit.6.2.D1.wahoo12\_wtuvell-  
May23/sbin/crosscc/spu10/sysroot/lib/libpthread.so.0

#1 0x081a471e in \_raise\_dfl (sig=6) at  
/home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1466

#2 0x081a47d9 in \_crash\_handler (sig=6, si=0x878b2ac, ignore=0x878b32c) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1570

#3 <signal handler called>

#4 0x452d45a6 in raise () from /nz/kit.6.2.D1.wahoo12\_wtuvell-May23/sbin/cross-cc/spu10/sysroot/lib/libc.so.6

#5 0x452d7d18 in abort () from /nz/kit.6.2.D1.wahoo12\_wtuvell-May23/sbin/cross-cc/spu10/sysroot/lib/libc.so.6

#6 0x082eec1f in CError\_abort (err=@0xffb452cc) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/error.cpp:459

#7 0x082eec96 in CError\_AssertFailed (line=699, src=0x8462728 "/home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp") at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/error.cpp:525

#8 0x0825e420 in CHash::HashTableFill\_Join (this=0x880f7690, numCols=1, colNums=0x73341b64, khash=true, simpleHasher=eInt32, stats=@0xffb47620) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp:699

#9 0x0825c95c in CHash::HashTableFill (this=0x880f7690) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp:492

#10 0x0825cb78 in CHash::CHash (this=0x880f7690, pPlan=0x7333eba4, pRS=0x7334a8b0, pHJ=0x73341d74, numCols=1, colNums=0x73341b64, hashStyle=eHSbasic2D, estimatedNumRecords=2639569, khash=true) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp:391

#11 0x0829f260 in DHJLooper (lastCall=<value optimized out>, pNode=<value optimized out>, pPlan=<value optimized out>) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/xplan.cpp:653

#12 CPlan::GenCodeLoop (this=0x7333eba4, lastCall=true) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/xplan.cpp:984

#13 0x0829faba in CPlan::ExecGenCode (this=0x7333eba4, pBufStart=0x0, pBufFence=0x0, lastCall=false, forHost=false, logIt=true) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/xplan.cpp:1069

#14 0x0814979c in \_do\_gened\_code\_final\_call (plan=<value optimized out>) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spuevent.cpp:2068

#15 \_handle\_scan\_io\_complete (dj=<value optimized out>, plan=<value optimized out>) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spuevent.cpp:2128

#16 spueventEmuJobSink::processJob (this=0xffb47c04, job=0x878fcdc) at

/home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spuevent.cpp:1798

#17 0x08256b30 in emuJobFilter::filterPassthrough (source=@0xffb479c8, sink=@0xffb47c04, timingAccumJob=0xffb47ad4, noEmu=false) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/comp/emu/emuJobFilter.cpp:141

#18 0x08256f4e in emuJobFilter::filterAsynchronous (source=@0xffb47be8, sink=@0xffb47c04, mode=emuModeProcess64, threadCount=6, timingAccumJob=0xffb47ad4) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/comp/emu/emuJobFilter.cpp:97

#19 0x0825711f in emuJobFilter::filter (source=@0xffb47be8, sink=@0xffb47c04, passthrough=false, timingAccumJob=0xffb47ad4) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/comp/emu/emuJobFilter.cpp:84

#20 0x081485db in SpuProcessPlanEvents (plan=0x7333eba4) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spuevent.cpp:1966

#21 0x081a5cc0 in \_invoke\_entry (arg=<value optimized out>, entry=<value optimized out>)

---Type <return> to continue, or q <return> to quit---

at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:737

#22 \_childproc (jobtask=0x78b4f924, priority=10, time\_slice=10, pipe=@0xffb47f10) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:812

#23 0x081a5e9f in \_\_create\_task (jobtask=0x78b4f924, priority=10, time\_slice=10, pipe=@0xffb47f10) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spu-jobtask.cpp:851

#24 0x081a5fe9 in \_create\_tasks (req=0x70563ba4, priority=10, time\_slice=10) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:916

#25 0x081a62e1 in \_create\_job\_tasks (req=0x70563ba4) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1294

#26 0x081a73c3 in nzprocmgr\_main () at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spujobtask.cpp:1753

#27 0x08141dce in NzSpuMain () at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spumain.cpp:1054

#28 0x08142436 in main (argc=Cannot access memory at address 0x7b90 ) at /home/workspaces/wtuvell/wahoo12/main/src/nde/spu/spumain.cpp:1064

(gdb) frame 8

#8 0x0825e420 in CHash::HashTableFill\_Join (this=0x880f7690, numCols=1, col-Nums=0x73341b64, khash=true, simpleHasher=eInt32, stats=@0xffb47620) at /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp:699

699 /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp: No such file or directory in /home/workspaces/wtuvell/wahoo12/main/src/nde/sys/hashrecs.cpp

(gdb) p totalRecordBytes

\$1 = 73834572

(gdb) p totalRecordBytes & 3

\$2 = 0

(gdb) p numRecords

\$3 = 2636949

(gdb) p \*m\_dataStore

\$4 = {m\_pRecords = 0x0, m\_NumRecords = 2639569, m\_allocatedRecSlots = 0, m\_hash = 0x0, m\_ModelTable = 0x7334a5b8, m\_pageCnt = 565, m\_rsResidentPages = 565, m\_rsDirtyPages = 0, m\_rsMapPosLo = 0, m\_rsMapPosHi = 565, m\_rsBcastRsPgacct = 0x0, m\_storeType = CRecordStore::eRStrasientTable, m\_currentPolicy = CRecordStore::eSPwindowView, m\_pageBytes = 131072, m\_pageCapacity = 131072, m\_recArrOwner = 0, m\_fileOwner = 5481494467300064144, m\_extendPageNum = 565, m\_firstSeqNum = 0, m\_lastSeqNum = 564, m\_RSdsid = 7, m\_firstWindow = 0x87967b8, m\_lastWindow = 0x8798ae8, m\_loopCursor = 0x0, m\_file = 0x8796130, m\_DPIBList = 0x8798630, m\_unusedDPI = 0x8798ad8, m\_storeLink = {next = 0x8080d6b8, prev = 0x73346f14}, m\_storeLinkPlan = 0x7333eba4, m\_spuImplicitClose = 5481494467300064144}

(gdb)

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Daniel Dietterich, Fritz Knabe, Jerrold Richard Title  
Bcc: Daniel Feldman  
Date: 06/01/2011 08:13 AM  
Subject: Re: Wahoobox is down

Yes, the 1GB succeeded, yesterday afternoon. So I kicked off a 100GB create be-



fore I left the office, and it built in a "good" time (4 hrs), then I kicked off a test just before I went to bed. I checked this morning, and it hung on query019, then I saw this note.

That was Plan A, and there's a Plan B in the wings: Rich's build. So I'm going to turn to that now (i.e., I don't think we can afford to debug this now, sorry).

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Fritz Knabe, Devesh Agrawal, Daniel Dietterich, Jeffrey Keller  
Bcc: Daniel Feldman, Steve Lubars  
Date: 06/01/2011 09:20 AM  
Subject: Way-back problem

Rich, I'm hitting a problem in your way-back build (debug). See wahoobox:/nz/Bar5.1/initialize/scripts. The "scriptA.sh" in that directory is the DB-build command I used, specifically "nohup ./scriptA.sh &>nohup.out &". Just prior to issuing this command, I did a "~nz/resart\_wahoo.sh -S 1 -D 8 -R". That crashed the first time I did it, so I did a mm001 power cycle and tried a second time, just in case that might clear things up. Then I tried a second time. That brings us up to date.

Following is the tail end of the nohup.out. Was there a hardware setup change with virtualization that could cause this?

DBGEN2 Population Generator (Version 1.0.0g-1)

Copyright Transaction Processing Performance Council (TPC) 2001 - 2008

(bootsvr) Info: SPU spu0101 (HWID 1012) (10.0.2.118): core dump request ACCEPTED

(bootsvr) Internal Error: Process 'spu10' pid 4211 exited (core dump) signal 11 (Segmentation fault)

(sysmgr) Info: NZ-01500: system state change from 'Online' to 'Pausing Now'

(sysmgr) Info: [spu hwid=1012 sn="e41f13e327b0" SPA=1 Parent=1002 Position=1 spuName= spu0101] got a heartbeat after 16 seconds, last heartbeat: Wed Jun 1 09:08:42 2011

(bnrmgr) Info: changing bnr mgr state to 'pausingNow', stopping current sessions

(bnrmgr) Info: bnr mgr completed transition to state 'pausedNow', version '6'

Error: Transaction rolled back due to restart or failover

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Error: Load Failed, records not inserted.

Error: Transaction rolled back due to restart or failover

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Load session of table 'CUSTOMER\_DEMOGRAPHICS' completed successfully

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Error: Load Failed, records not inserted.

Error: Transaction rolled back due to restart or failover

Error: Transaction rolled back due to restart or failover

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Error: Transaction rolled back due to restart or failover

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Error: Transaction rolled back due to restart or failover

See /tmp/wahoobox/wahoobox/tpcds100b/CUSTOMER\_DEMOGRAPHICS.TPCD-S100.nzlog file

Error: Load Failed, records not inserted.

Error: Load Failed, records not inserted.

Error: Load Failed, records not inserted.

Error: Load Failed, records not inserted.

Error: Load Failed, records not inserted.

2011-06-01 09:09:05|FATAL||PID:5440 accessing host wahoobox finished with status:256

(sysmgr) Info: NZ-01500: system state change from 'Pausing Now' to 'Going Pre-Online'

(fcommrtx) Info: fcommTransChanPoolMb 452 fcommPermChanPoolMb 3

(fcommrtx) Info: \_validateDevShmSize(): gShmMb 1540 commPoolsMb 455 qcLoadMb 17580 devShmMb 78116

(sysmgr) Info: NZ-01500: system state change from 'Going Pre-Online' to 'Resuming'

(alcapp) Info: Initialized Interface Settings: BurstSendInterval:10; BurstSendRetries:1; LongSendInterval:1000; LongSendRetries:-1; LongSendOnQhist:0; SocketReadyTimeout:2000.

(sysmgr) Info: NZ-01500: system state change from 'Resuming' to 'Online'

(bnrmgr) Info: system is online - enabling backup and restore sessions

■ From: Daniel Feldman  
To: Walter Tuvell  
Date: 06/01/2011 09:28 AM  
Subject: Re: Way-back problem

I'm sure none of this is new to you but think it is worth discussing:

Your email below prompted me to think about debug builds. In general, it is not a good practice to rely on results from debug builds. This is because debug builds often include code that adds significant incremental processing for the system under test. The issue is complicated even more if a baseline is established using non-debug (production) code and the system being assessed against the baseline is run using debug code. Not that debug v debug assures commensurability, but production v debug certainly strongly suggests its absence.

How do you assess the debug code v production code issue in the context of this particular testing task?

■ From: Walter Tuvell  
To: Daniel Feldman  
Date: 06/01/2011 09:32 AM  
Subject: Re: Way-back problem

This was discussed in Camb yesterday. We're using debug first to try to get the damn thing through a test, then switch to turbo for the numbers.

- From: Walter Tuvell  
To: netezza-cambridge  
Date: 06/01/2011 09:37 AM  
Subject: Standup today

I have a conflicting mtg during standup, but I've been in contact with Fritz & others about my status, so they can probably combine to fill in the details, but the short story is that attempts to run TPC-DS@100 on WahooBox are still failing (multiple modes).

- From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Devesh Agrawal, Daniel Dietterich, Fritz Knabe, Jeffrey Keller ►Forwarded to Dan Feldman, 06/01/2011 10:04 AM.◄  
Date: 06/01/2011 10:02 AM  
Subject: Re: Wahoobox is down

The core dump stack trace (below) shows it's crashing in pm::DsDesc::addPrecommittedTXid() (trying to insert a txid into m\_precommittedTXids).

This appears to be code Jeff added for prototype garbage collection on Apr 14, based on Accurev histories. My snapshot was for Apr 27.

So we do have some of the prototype GC in the "wayback" snapshot, and it's causing problems.

Not sure what to suggest. We could back up even further, but I'm not sure if that's a productive path to be going down.

►Screenshot omitted here; irrelevant to this Complaint.◄

- From: Walter Tuvell  
To: Fritz Knabe  
Bcc: Daniel Feldman  
Date: 06/01/2011 10:09 AM  
Subject: DITW

FYI, I'm now dead-in-the-water.

- From: Jerrold Richard Title  
To: Jerrold Richard Title  
Cc: Walter Tuvell, Devesh Agrawal, Daniel Dietterich, Fritz Knabe, Jeffrey Keller  
Date: 06/01/2011 10:32 AM  
Subject: Wayback machine

If we're going to back up further, the previous point where we synced up with int-trunk was 3/11/2011.

Our recent merge history looks like this:

```

dev-wahoo -----+-----+-----+---> arrow of time
3/11             4/27             (now)
merge           merge           merge
w/trunk         w/trunk         w/trunk
in merge hell:) (Dan)           (Rich)           (Devesh)

```

Backing up to points where we were not sync'ed up with int-trunk is problematic. For instance, if I were to put a timebasis of 4/1 on my wayback stream, it'd "see" things in int-trunk that are overlapped with what it sees in dev-wahoo on 4/1.

Should I try making a 3/11 "wayback" build, or is that too far back?

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal, Daniel Dietterich, Fritz Knabe, Jeffrey Keller  
Bcc: Daniel Feldman  
Date: 06/01/2011 10:36 AM  
Subject: Re: Wayback machine

No need, I've got a Mar 30 build already, in fact it's the one I used to make that pic on our whiteboard.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal, Daniel Dietterich, Fritz Knabe, Jeffrey Keller  
Bcc: Daniel Feldman  
Subject: Re: Wayback machine

As I understand it, I'm going to try reviving my Mar 30 build, while Rich takes another trip down memory lane as a backup. I'll start on my part now.

■ From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Fritz Knabe, Daniel Dietterich, Jeffrey Keller, Devesh Agrawal  
Date: 06/01/2011 02:18 PM  
Subject: for your time travelling pleasure

In case you need it, March 11 2011 vintage dev-wahoo builds are available here:

</net/rtitlerhl/home/workspaces/rttitle/wahoo-wayback-3-11/main/debug>

/net/rtitlerhl/home/workspaces/rtile/wahoo-wayback-3-11/main/turbo

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Fritz Knabe, Daniel Dietterich, Jeffrey Keller, Devesh Agrawal  
Bcc: Daniel Feldman  
Date: 06/01/2011 03:18 PM  
Subject: Re: for your time travelling pleasure

Thanks Rich. I've been running my Mar 30 build, and it hasn't hung (yet :-), but it's unaccountably running a good deal slower than it did in Mar. So bad, in fact, that I'm going going to kill it forthwith and use yours. Fingers crossed, everyone ...

■ From: Walter Tuvell  
To: netezza-cambridge  
Date: 06/01/2011 03:59 PM  
Subject: From the frying pan into the fire

Well, this isn't going as well as could be expected. Rich, this is your Mar11 debug. Any ideas?

(bootsvr) Info: NZ-00022: --- program 'bootsvr' (18048) starting on host 'wahoo' ... ---

(bootsvr) Info: Responded to boot request from device [ip=10.0.2.118 SPA=1 Slot=1] Run Level = 5

(bootsvr) Info: Responded to boot request from device [ip=10.0.2.118 SPA=1 Slot=1] Run Level = 5

(bootsvr) Info: Responded to boot request from device [ip=10.0.2.118 SPA=1 Slot=1] Run Level = 5

(sysmgr) Info: 4 SCSI error codes read from /net/rtitlerhl/home/workspaces/rtile/wahoo-wayback-3-11/main/debug/sys/scsi\_error\_policy.cfg:

(sysmgr) Info: - 0x04 0x32 0x00

(sysmgr) Info: - 0x04 0x32 0x01

(sysmgr) Info: - 0x04 0x42 0x00

(sysmgr) Info: - 0x04 0x42 0x01

nzstart: Warning: system did not go online after 300 seconds.



[nz@wahoobox nz]\$

[nz@wahoobox nz]\$

nz@wahoobox nz]\$ (bootsvr) Info: Responded to boot request from device  
[ip=10.0.2.118 SPA=1 Slot=1] Run Level = 3

(bootsvr) Info: Responded to boot request from device [ip=10.0.2.118 SPA=1  
Slot=1] Run Level = 5

(sysmgr) Warning: Recieved initCall from spu with different ip's: ipAddr:  
10.0.2.118, infoIpAddr: 10.0.2.117

(sysmgr) Info: NZ-01545: added a new spu [spu hwid=1003 sn="e41f13e327b0"  
SPA=1 Parent=1002 Position=1 spuName= spu0101] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1004 sn="WD-  
WXL508081242" SPA=1 Parent=1003 Position=1] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1005 sn="WD-  
WXL508080038" SPA=1 Parent=1003 Position=2] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1006 sn="WD-  
WXL508054390" SPA=1 Parent=1003 Position=3] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1007 sn="WD-  
WXL508055392" SPA=1 Parent=1003 Position=4] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1008 sn="WD-  
WXL508055394" SPA=1 Parent=1003 Position=5] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1009 sn="WD-  
WXL508055395" SPA=1 Parent=1003 Position=6] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1010 sn="WD-  
WXL508055396" SPA=1 Parent=1003 Position=7] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1011 sn="WD-  
WXL508055397" SPA=1 Parent=1003 Position=8] to the system

(sysmgr) Info: NZ-01545: added a new disk [disk hwid=1012 sn="WD-  
WXL508079993" SPA=1 Parent=1003 Position=9] to the system

(sysmgr) Info: Sending InitReply to hwid=1003

(sysmgr) Error: The system is DOWN because disk with address  
disk.spa.1.slot.1.lp.1 was not found by spu spu.spa.1.slot.1none.

(sysmgr) Info: NZ-01500: system state change from 'Discovering' to 'Down'

[nz@wahoobox nz]\$

[nz@wahoobox nz]\$ nssystem

(clientmgr) Info: admin: login successful

System state is 'Down'.

Expected state is 'Discovered', version 2.

[nz@wahoobox nz]\$

■ From: Huamin Chen  
To: Walter Tuvell  
Cc: netezza-cambridge  
Date: 06/01/2011 04:01 PM  
Subject: Re: From the frying pan into the fire

looks like a bad topology file.

■ From: Jerrold Richard Title  
To: Huamin Chen  
Cc: Walter Tuvell, netezza-cambridge  
Date: 06/01/2011 04:18 PM  
Subject: Re: From the frying pan into the fire

Right. I think the format of the topology file changed between Mar 11 and now. The older code base expects .pp. instead of .lp. in those disk addresses (Devesh just told me this). Devesh and I just edited his restart\_wahoo.sh script to generate the older format, and I'll re-run it right now, and we'll see if wahoobox comes up.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Huamin Chen, netezza-cambridge  
Date: 06/01/2011 04:20 PM  
Subject: Re: From the frying pan into the fire

Thank you, gents.

■ From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Huamin Chen, netezza-cambridge

Date: 06/01/2011 04:52 PM  
Subject: Re: From the frying pan into the fire

Sorry, we're still having some problems here bringing that Mar 11 kit up on wahoobox.

I think stuff has changed between Mar 11 to now in ways that confuse the bring-up scripts. Devesh is poking around.

It did come up on a Dwarf, but that doesn't help you. :-(

■ From: Jerrold Richard Title  
To: Jerrold Richard Title  
Cc: Walter Tuvell, Huamin Chen, netezza-cambridge  
Date: 06/01/2011 05:49 PM  
Subject: Re: From the frying pan into the fire

OK, we have a theory. Basically, bad choice of timebasis. (Turns out Mar 11 was not a good vintage for wahoobox-flavored wine :-).

I had set the timebasis to March 11 6pm, chosen because it was just after our March 11 merge-with-all-the-dolphin-stuff.

But, it looks like wahoobox startup was actually busted on that date. Devesh promoted a set of changes on March 15 with the checkin comment "wahoobox startup works now". So, I'm going to move my timebasis forward to March 15 (after Devesh's promotion) and rebuild and try again to bring up wahoobox.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Jerrold Richard Title, Huamin Chen, netezza-cambridge  
Date: 06/01/2011 06:03 PM  
Subject: Re: From the frying pan into the fire

FYI, I'll be out of the office for the next 2 days, partially working, partially not, but in any event I'll be monitoring this situation and acting upon it as events dictate.

Just in case: 781--475-7254 (c), 781-944-3617 (h).

■ From: Jerrold Richard Title  
To: Walter Tuvell  
Cc: Huamin Chen, netezza-cambridge  
Date: 06/01/2011 06:34 PM  
Subject: Re: From the frying pan into the fire

Still no luck.

Going back that far in time appears to be not so easy. It's not clear if we're failing to capture a stable point, or if Accurev is failing to give us an accurate worldview as of the specified time, or some other inconsistency.

It doesn't feel productive to be debugging a 2-month-old snapshot. Let's figure out tomorrow what makes sense as a plan going forward. Can we call you at one of those numbers during the day tomorrow?

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Huamin Chen, netezza-cambridge  
Date: 06/01/2011 08:26 PM  
Subject: Re: From the frying pan into the fire

Yes, try cell-phone first (wife uses landline), and alert me by email if you can. I was supposed to be working on my deck, but the carpenter bailed (got sick) at the last minute. However, I'll still be working-at-home, because I have an ulterior motive: I'm having a colonoscopy on Fri, hence doing that crazy diet thing.

FWIW: That Mar 30 build I tried, even though it was indeed the one I used successfully previously, but was weirdly too slow this time (and seemed nearly hung when I killed it), was debug, and I don't still have the source so I can't generate a turbo. When it rains, it pours.

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Jerrold Richard Title, Huamin Chen, netezza-cambridge  
Date: 06/02/2011 09:03 AM  
Subject: Re: From the frying pan into the fire

Walt, based on what Rich was able to see this morning, it looks like the Mar 30 build is running. Is it still showing problems? Unless it is proceeding well, let's do this:

1. Walt: Kill the test. We'll give up on old builds.
2. Steve or anybody: Reconfigure wahoobox to two virtual machines.
3. Rich, Devesh, et al: Proceed with the merge and with debugging forward on the tests. Use wahoobox for testing. Our goal is to get Wahoo running again as quickly as we can so that Walt can performance test on the current state of the system.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Jerrold Richard Title, Huamin Chen, netezza-cambridge  
Date: 06/02/2011 09:48 AM  
Subject: Re: From the frying pan into the fire

Sorry, I was out getting a blood test (medical catch-up day, sigh).

I started a Mar30 run last night, but it hung at query019. Seeing that this morning, killed it, removed query019 from the list, then restarted. Now that I'm back, I'll take a look.

■ From: Walter Tuvell  
To: Fritz Knabe, Jerrold Richard Title, Huamin Chen, netezza-cambridge  
Bcc: Daniel Feldman  
Date: 06/02/2011 09:55 AM  
Subject: Re: From the frying pan into the fire

It was hung, so I killed it. Steve, you can do the reconfig now.

I thought the slowness of the Mar30 build yesterday may have been due to the fact that I wasn't generating stats, though I've done it both ways in the past and seen little measurable difference (at least on Atomics & TPC-DS). So this last test did gen stats, but it hung doing that (never got to the query019 it hung on last night). Too weird.

■ From: Devesh Agrawal  
To: netezza-cambridge ►*Forwarded to Dan Feldman, 10:11 AM.*◄  
Date: 06/02/2011 10:09 AM  
Subject: dev-wahoo status ... still broken.

Hi Folks,

I would like to update you on the status of getting dev-wahoo back.

I am still chasing a data corruption kind of a bug on a real machine. My wild guess is that something to do with async writing is broken with the new code we have taken. I see this problem occurs on sqlsmoke (emp table tests itself) but yet it does not manifest when I load and scan GB's of tpcds data. The likely culprits that may have caused this problem could be (wild guesses again) : HostTXMgr changes or Andy's message deferral thingy in spuevent.

Still groping .... :( When this works I would like to test this on wahoobox too. Unit tests and vspu works atleast.

■ ►*Here occurred the “next steps” email chain (Appendix F).*◄

■ From: Devesh Agrawal  
To: netezza-cambridge  
Date: 06/03/2011 09:03 AM  
Subject: Wahoo is safe again

Hi Folks,

It is safe to update from dev-wahoo now. I removed the async writer (stubbed it out) since it was leading to some data corruption with Andy's (probably) recent changes. In the meantime, I will figure out what's up it ...

Thanks.

PS: Jeff, I commented out one of your gcTsLogSnapshot tests since that was failing and I couldn't figure out why. I also removed the activate/de-activate tests since they were probably exercising some sort of fcomm bug --- spu fcomm used to crash with it.

■ From: Jerrold Richard Title  
To: Devesh Agrawal  
Cc: netezza-cambridge  
Date: 06/03/2011 09:31 AM  
Subject: Re: Wahoo is safe again

Great! Thanks for all the hard work on this merge.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: Devesh Agrawal, netezza-cambridge  
Date: 06/03/2011 02:53 PM  
Subject: Re: Wahoo is safe again

Yes, thanks Devesh, Merge Hell isn't much fun under the best of circumstances, and especially bad when people are breathing down your neck. FWIW, my experiment from earlier this morning (on Atomics) hung too, so I'll upgrade. Talk to you all next week.

■ From: Steve Lubars  
To: Walter Tuvell  
Cc: Jerrold Richard Title, Devesh Agrawal  
Date: 06/03/2011 04:12 PM  
Subject: Re: Wahoo is safe again



Hi Walt -

Were you planning to run another test this weekend with the upgraded build?

■ From: Walter Tuvell  
To: Steve Lubars  
Cc: Jerrold Richard Title, Devesh Agrawal, Fritz Knabe  
Date: 06/03/2011 04:29 PM  
Subject: Re: Wahoo is safe again

Yes, I've already updated and am now building. I'll use Wahoo1, so Wahoo2 is free for others.

■ From: Steve Lubars  
To: Walter Tuvell  
Date: 06/03/2011 04:33 PM  
Subject: Re: Wahoo is safe again

Do you anticipate that NMON will give us network stats similar to the CPU, EMU, and IO status we got from the last run? I'm worried that we'll find out everything is good to go on Monday, but I won't be ready to collect data; Fritz is suggesting I do a practice run of Atomics (somewhere) just to work out the details of spreadsheet generation...

■ From: Walter Tuvell  
To: Steve Lubars  
Date: 06/03/2011 05:19 PM  
Subject: Re: Wahoo is safe again

Well, nmon (in my modified version, now called nzmon) work very well on host, but I haven't tried it on SPU yet. My previous data capture/report tools work only on SPU, not host. Furthermore, I don't have any reporting stuff whipped up for the nzmon format yet (the previous tools work on a different format). So you're right, there is a gap here.

How do you want to handle this? You can use skimmer (s1-8) if you want, but I'm booked this weekend, except for the few hours I'll be able to steal to try to get the new kit working on Wahoo.

■ From: Steve Lubars  
To: Walter Tuvell  
Date: 06/03/2011 05:24 PM  
Subject: Re: Wahoo is safe again

I won't have much time to dedicate this weekend either, but in case I do I want to spend it effectively... some choices:

- 1) I can try adding network IO to your previous data capture/report tools (for the SPU only)
- 2) I can try getting previous data capture/report tools to work on the Host
- 3) I can try running nzmon on the SPU
- 4) I can try getting nzmon to output in Excel-friendly format

My inclination is to try (3), but I'm open to any of them - can you provide pointers to the nzmon stuff (and/or the previous tool)?

■ From: Steve Lubars  
To: Walter Tuvell  
Date: 06/03/2011 05:38 PM  
Subject: Re: Wahoo is safe again

Just had a conversation with Devesh and Fritz - they feel the highest priority task is to verify that Wahoo (with the latest round of changes) performs as well as it did before, even if it hangs or dies before all the tests have been run. To that end, can you capture data (on the SPU, using the previous tools) when you run your tests?

■ From: Walter Tuvell  
To: Steve Lubars  
Date: 06/03/2011 05:38 PM  
Subject: Re: Wahoo is safe again

In general we should start migrating toward nzmon, because my previous way of doing it was a one-off hack, even though it works fine, until I could figure out how to do it better (the usual "one to throw away" syndrome).

To the extent you have some free time this weekend, you should probably start off with something easy, say #4, just to get an idea of what we're talking about. Just Google nmon. There are already 2 tools for Excel-ization (as well as 1 for RRDTool), which I haven't tried, and it would be very useful if you could get that going. If you get it working for vanilla nmon, I can get it working for nzmon (all this is open-source). All you need to do is get it working on any convenient Linux box, no need for Wahoo or NPS of any flavor.

As for #3, either it'll be easy or hard. If easy, I can do it, and then we can decide if we want to use it, depending on how #4 comes out. If hard, it'll have to wait until after this exercise.

How does that sound?

■ From: Walter Tuvell  
To: Steve Lubars  
Date: 06/03/2011 05:40 PM  
Subject: Re: Wahoo is safe again

OK, thx.

■ From: Steve Lubars  
To: Walter Tuvell  
Date: 06/03/2011 05:56 PM  
Subject: Re: Wahoo is safe again

Thanks Walt.

I'll play with nmon output this weekend (wife/kids permitting).

►*Above thread "Re: Wahoo is safe again" forwarded to Dan Feldman, 06/03/2011 07:07 PM.*◄

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Fritz Knabe, Steve Lubars, Jerrold Richard Title  
Date: 06/03/2011 07:29 PM  
Subject: Not booting

Devesh the build you pointed me to today isn't booting. I've tried both debug and turbo. This is wahoo1, if you'd care to take a look.

■ From: Devesh Agrawal  
To: Walter Tuvell  
Cc: Fritz Knabe, Jerrold Richard Title, Steve Lubars  
Date: 06/03/2011 10:49 PM  
Subject: Re: Not booting

I had to power cycle the spu1 via the amm. Sometimes that helps as a way to force a hard restart. Now wahoo1 is up with my turbo kit.

Btw, wahoo2 is also up with the same kit. So both are yours.

PS: Btw, a debug build ofcourse won't work because I only built the turbo version :-).

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Fritz Knabe, Jerrold Richard Title, Steve Lubars  
Bcc: Daniel Feldman  
Subject: Re: Not booting

Date: 06/04/2011 10:26 AM

Dammit, I should have thought of power cycling! The reason I didn't is that the symptom this time was different from what I've seen before that power cycling solved. Live and learn. Thx.

Anyway I was able to build the 100G DB (with parallelism, 68 min, yay!), and then tried running Atomics, but it hung (multi-hours) on join\_multi\_hash:

```
-- JOIN-MULTI-HASH Compute-intensive, co-located, multi-hash star join.
SELECT AVG(cs1.cs_ext_sales_price)
FROM catalog_sales cs1
  INNER JOIN catalog_sales cs2 ON (
    cs1.cs_order_number = cs2.cs_order_number
    AND cs1.cs_item_sk = cs2.cs_item_sk
    AND cs1.cs_bill_customer_sk != cs2.cs_ship_customer_sk
  )
  INNER JOIN catalog_returns ON (
    cs1.cs_order_number = cr_order_number
    AND cs1.cs_item_sk=cr_item_sk
  )
  INNER JOIN date_dim ON (cr_returned_date_sk = d_date_sk)
WHERE d_date BETWEEN '1999-01-01' and '1999-02-01'
;
```

On the other hand, I looked at the partial Atomics result-summary file, and noticed that the cross\_join test took nearly 1 hr. Recalling that we'd earlier in the week seen a hang in cross\_join, I wondered whether there might be a relationship between these two tests (maybe the word "JOIN"?):

```
SELECT COUNT(*)
FROM web_returns c1
CROSS JOIN date_dim
--WHERE d_date_sk<2415022+1000;
--WHERE d_date_sk<2415022+500;
WHERE d_date_sk<2415522;
```

After the Atomics hang, one thing I could have done was to remove the join\_multi\_hash and restart Atomics. But instead I decided to try TPC-DS (which is our first priority, though both are important). That's now running, currently at query031, but it's been there a long time, nearly 30 min. I'll leave it running, but I wouldn't be surprised to see it hang there. If anybody wants to check, you can look on wahoo1 at /nz/Bar5.1/scripts/nohup.out.

BTW, I'm using a collection of settings Devesh and I worked out. No surprises, all

of these have been mentioned before, but for the curious I'm including an attachment recording them.

As for the turbo vs. debug problem Devesh mentions, I guess I was misled when he told me in IM that he'd only built this kit debug to date. I should have realized the empty set of builds satisfies his statement! :-)

Finally, I want to report a build bug in dev-wahoo. I can't build "image" (the target) now, though I've been able to build it in the past. This is important for me, because that's the way I build (because I don't use a workstation [nor do I want to, done enough of that in my life], so need to transfer the software to the machine). Hence, I'm now running with Devesh's build (turbo, which isn't ideal, because we want to be able to debug problems as they occur, then switch to turbo for number-gathering, as we've discussed). So somebody will have to look into this, unless they want to be bothered to do all my building for me. :-)

See you on Mon.

■ From: Devesh Agrawal  
To: Walter Tuvell  
Cc: Fritz Knabe, Jerrold Richard Title, Steve Lubars  
Date: 06/04/2011 12:07 PM  
Subject: Re: Not booting

I didn't see any core files and the spu didn't undergo a state change. It is hung on the same WLM'ish bug that Dan was investigating a week or so ago. Basically we enter this hammer mode thingy on the host, where it just `_waits_` indefinitely. (The spu is all done).

I am not sure if perfbar is doing any GRA related hacks which could be inducing this or if there is any way to turn off this WLM/GRA crap all-together and let the queries go at full throttle. You might want to (or I will be happy to do that) circle with Dan Barrett and Gordon about this.

I would suggest killing it for now.

PS: I am using wahoo2.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Fritz Knabe, Jerrold Richard Title, Steve Lubars  
Bcc: Daniel Feldman  
Date: 06/04/2011 06:37 PM  
Subject: Re: Not booting

Thanks for inspecting, Devesh. It is indeed still hung on query031 I'll kill it, and

try building a 1000G DB, just for the hell of it (wouldn't want all those good cycles going to waste).

Since I don't really know how what is required to sidestep WLM/GRA-ish problems, it would be great if you'd talk to DanB & Gordon. That will give me time to work some more on stats gathering/reporting (nzmon, Steve, etc.).

See you-all on Mon.

■ From: Devesh Agrawal  
To: netezza-cambridge ►*Forwarded to Dan Feldman, 06/06/2011 05:53 AM.*◄  
Date: 06/05/2011 06:34 PM  
Subject: Lessons learnt from my multi-slice data corruption bug -- aka a summary of DsIndex, DevNo, DataSliceId.

DevNo, DsIndex and DataSliceId: The fore-fathers of NPS's storage layer were probably lawyers in their past life to come up with these three concepts which are so confusingly similar (yet different !) and ofcourse without any explicit documentation.

I spent a good part of the past week chasing a bug which had to do with my jumbling these concepts. So now that I have learnt the hard way, let me write this down for posterity :-).

- Data Slice Id is a system wide (not a blade wide) id for a data slice. The data slices are stored in the NzState->dataSliceTbl and are numbered 1..N. A Spu can be assigned arbitrary data slices.

- Data Slice Index is a bit confusing. There are actually two data slice index's. One is mt::dsidIndex that maps a data slice id to its index. This is basically the index of the array mt::mt\_dsids that matches the given data slice id. BUT DPM and PM don't use this.

What DPM and PM do is: They have their own array of data-slice's. In pm's case this is the m\_dsDesc array in DsGlobal. Entries in this array used to be assigned in order (the mounting was done one at a time for each slice). That changed with this merge and now the mounting of slices is done by concurrent threads. So -- the data slice index's can be arbitrary. (ie have no relation to the data slice id).

What does the DsIndex do ? Basically we store the in-memory metadata about a dataslice A in its dsIndex(A). In our case, the DsDesc's are stored in the array m\_dsDesc in DsGlobal in DsIndex order. Probably back in the day when they were inventing this, hash table's didn't exist :-).

- Okay so now what is this Device Number ?. Device numbers range from 0 to NumSlicesOnSpu - 1. They basically indicate the nzds partition that the storage layer is talking to. So when we write to data slice A we are actually writing to device DevNumber(A) and ditto when we scan from A.

So my bug was : (and the code review never caught this !) I confused the DsIndex with DevNo. So I was sending my writes to the wrong partition effectively which was leading to all this corruption business. The reason this was hidden before the time-basis removal was: that the mounting was single threaded, so dsIndex == devNumber. Not so any longer !.

AAAAAAAAAAAAARRRRRRGH ..... and I was doubting the sanctity of everything under me like the linux block layer etc. (How can my own code be buggy !).

PS: I will promote the fix tomorrow after cleaning up a lot of cruft that I had to cast to zero in on this.

■ From: Devesh Agrawal  
To: Daniel Barrett, Gordon Booman  
Date: 06/04/2011 08:53 PM  
Subject: Turning off WLM/GRA completely

Hi Folks,

We are having two GRA related problems when we are running tpcds queries on our frankenstein wahoo prototype. And I think they are because we (or the WLM team) haven't gotten around to tweaking any WLM stuff for Wahoo yet. The two problems are:

(a) DanD has already discussed this one with you. For some reason we enter into this hammer-mode thingy in the Interaction Wrap up (when the spu is done) and the dbos event process just waits over there indefinitely and the whole thing has to be killed. This usually happens when a temp-table (not transient table) is being downloaded to on the spu. And this is a random event -- does not happen reliably all the time. So it is hard to pinpoint to a query and say that this will be triggered for sure. I would also like to point out that we are running a single query stream -- nothing else is running on the system.

(b) When we look at the GRA virtual tables, we find that the perfbar created user has been placed in a bucket that has only 1% resources !. This is clearly a bug but I am not sure whether this is a perfbar bug wherein it is calculating a wrong percentage (1%) for us. Or is there something wrong in the GRA code for Wahoo ?.

In any case, what I would like to do is to turn off all such GRA related stuff to achieve no artificial query slowdown whatsoever. Is there a way to do this ?

■ From: Gordon Booman  
To: Devesh Agrawal  
Cc: Daniel Dietterich, netezza-wlm-team (netezza-wlm-team@wwpdl.vnet.ibm.com)  
<netezza-wlm-team@wwpdl.vnet.ibm.com>



Date: 06/05/2011 10:11 PM  
Subject: Re: Turning off WLM/GRA completely

Hi Devesh,

Well, this is embarrassing. I'm sorry you're running into these problems.

Generally speaking, we don't expect any specific changes for Wahoo. There's no Wahoo mode or anything. We are looking at better ways to control very fast snippets; that's happening as part of Nighthawk. It's possible that we will need to do more for Wahoo, but right now, we don't have anything specific planned.

Regarding the first problem, hammer mode. As you said, Dan mentioned this. We added an assert for what we think you are seeing. That went in to trunk last week. I would love to see this happen so we can figure it out. The assert is our way of making that happen. :-)

But if you need to avoid this altogether, there is a registry setting you can use to turn off greed. Set `host.schedFlags=16`.

Regarding the second problem, would be happy to take a look. A 1% setting is not necessarily a bug. It can be a valid setting. Most devs run as admin, which is a special group and gets at least half of the system resource. I think most perfbar tests just run as admin. In the perfbar muqry test, different users are set up with different minimums. I don't recall the exact settings. I think it has changed. If you tell me the test (user, perfbar config & machine) you're running, I can take a look. Or if you run it and it gets in this state, I can take a look. That might be easier.

It is possible to turn off GRA, but then you can't run the entire perfbar since it requires GRA.

Anyway, to turn off GRA:

```
host.schedAllowGKandGRA=no # default is no
```

```
host.schedGRAEnabled=no # default is yes, "no" means that GK is on instead of GRA
```

This odd combination is for historic reasons.

■ From: Devesh Agrawal  
To: Fritz Knabe, Walter Tuvell  
Date: 06/06/2011 10:33 AM  
Subject: Fw: Turning off WLM/GRA completely

So we can either try to turn off GRA completely by setting those registry variables and then risk that perfbar may not run completely (although I am not sure whether our particular test suite will complain). If that does not help we would

need to get the latest code from trunk (re-remove the time-basis and do the whole dance again).

■ From: Devesh Agrawal  
To: Walter Tuvell, Jerrold Richard Title, Jeffrey Keller  
Date: 06/06/2011 03:25 PM  
Subject: dev-wahoo is back in good shape now

Sqlsmoke passes on vspu and non-wahoo (so does bar). Image and tarimage target works.

Walt, please build your own kit now. I will recycle my ws2 workspace for something else.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Jeffrey Keller, Jerrold Richard Title  
Date: 06/06/2011 04:22 PM  
Subject: Re: dev-wahoo is back in good shape now

OK, I've upgraded AccuRev, updated the workspace, and will build this evening. Gotta leave now.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Jeffrey Keller, Jerrold Richard Title, Fritz Knabe  
Date: 06/06/2011 08:38 PM  
Subject: Re: dev-wahoo is back in good shape now

Devesh, the "image" build succeeded, but it failed to unpack on Wahoo1. Here's the symptom:

►*Screenshot omitted; irrelevant to this Complaint.*◄

■ From: Walter Tuvell  
To: Devesh Agrawal, Jeffrey Keller, Jerrold Richard Title  
Cc: Fritz Knabe  
Date: 06/07/2011 08:36 AM  
Subject: Re: dev-wahoo is back in good shape now

Uh, forget that. I lied (either that or I'm an idiot, maybe both). I was trying to unpack on my laptop. Duh.

Noticing that flub this morning, I've succeeded in unpacking on Wahoo1. Now to see if it boots, builds 100G DB, runs TPC-DS test. Should know by stand-up time.

■ From: Devesh Agrawal  
To: Walter Tuvell  
Date: 06/07/2011 09:42 AM  
Subject: Re: dev-wahoo is back in good shape now

One more thing -- use the new perfMeasure.rb script -- it is checked into dev-wahoo. It has network monitoring also now.

■ From: Walter Tuvell  
To: Fritz Knabe  
Bcc: Daniel Feldman  
Date: 06/07/2011 03:12 PM  
Subject: Just when you thought things couldn't get worse ...

Devesh has just promoted his latest bugfix (I don't think you even knew about it, it's preventing NPS on Wahoo1 from booting, it keeps bouncing between Online and Pre-Online).

Anyway, so I went to get the fix from Accurev, and I can't login! It's not accepting my password.

I've contacted IT ...

■ From: Walter Tuvell  
To: helpdesk  
Date: 06/07/2011 03:08 PM  
Subject: Can't login

I can't login as wtuvell (on Netezza internal network). It's not accepting my password.

■ From: Help Desk  
To: Walter Tuvell  
Date: 06/07/2011 03:33 PM  
Subject: Re: [Request ID :##40723##] : Can't login

Your account is not showing locked... are you trying to get on wireless? please try netezza\wtuvell for the username. also check our selfservice utility to reset your password. <https://selfservice.netezza.com:9251> (use wtuvell for username there)

If you still can't get in I can reset your password for you.

Jay Griffin

■ From: Walter Tuvell  
To: helpdesk  
Date: 06/07/2011 03:47 PM  
Subject: Re: [Request ID :##40723##] : Can't login

Hi, Jay. I'm wired (in Camb), not wireless.

I don't need "netezza\wtuvell", because I'm just trying to "su wtuvell" on my Fedora laptop. This has always worked in the past. I checked I'm on the Netezza network, and I am (I can "cd ~wtuvell").

If I reset my password now, it'll be out-of-sync with my password everywhere else. [Yes, I know this is a weak security practice, but everyone will continue to do it until single-sign-on becomes a reality.]

Anyway, I just did try resetting my password via selfservice, and same problem: I can't login in to it. (I tried all the passwords I've used at Netezza/IBM, and none of them work.)

■ From: Help Desk  
To: Walter Tuvell  
Date: 06/07/2011 03:56 PM  
Subject: Re: [Request ID :##40723##] : Can't login

I have reset your password to

Netezza1234

you should be able to reset it to something more secure after three days. Please let me know if you continue to have issues...

■ From: Walter Tuvell  
To: helpdesk  
Date: 06/07/2011 04:00 PM  
Subject: Re: [Request ID :##40723##] : Can't login

OK, I'm logged in now, don't have any idea why it didn't work before. Surely I can't mistype my password 10+ times in a row can I????

Anyway, THANKS, as always

■ From: Walter Tuvell  
To: Steve Lubars  
Cc: Fritz Knabe, Daniel Feldman  
Date: 06/07/2011 08:36 PM  
Subject: Clock sync on NPS

Steve, I misspoke today when I told you I was synchronizing clocks to within 0.1 sec via my script suite. Actually, it's only to within ~1 sec (which is unfortunately at the very edge of coarse granularity we're able to use). The reason is that the stupid busybox "date" command only supports that granularity of time-setting. (Real Linux can do better.)

I've submitted a Jira ticket for NPS to provide true/NTP clock-sync within clusters, but I haven't seen any activity on it. IMHO, this is a high-priority item, due to the fine-granularity we're studying NPS within Wahoo. That's why I'm cc'ing Fritz & Dan, in the hope that if they agree they might help rally some votes to push the Jira through.

It's #SWS-65271.

(Sorry for my misspeak, there are only so many details I can keep in my head at once ...)

■ From: Walter Tuvell  
To: Fritz Knabe  
Date: 06/08/2011 08:23 AM  
Subject: Clock sync

```
#!/bin/bash
```

```
# This script runs on NPS hosts.
# It syncs host clock to NTP, then syncs blade clocks to host.
# Not as good as host serving NTP to blades, but not so bad as an occasional manual hack.
# Depends upon password-free root SSH access to blades.
```

```
# 0. Check NPS is running.
```

```
nzsystem &>/dev/null
```

```
[[ $? == 0 ]] || { echo "*** ERROR: You must boot NPS first (manually)."; exit 42; }
```

```
# 1. Sync host to NTP. This is straightforward, though perhaps a bit heavy-handed.
service ntpd status &>/dev/null # if NTPD is already running, we assume life is hunky-dory
```

```
if [[ $? -ne 0 ]]; then
```

```
    sudo ntpd -q # slam host clock manually (assumes /etc/ntp.conf is already configured, which it should be)
```

```

sudo service ntpd start &>/dev/null # start NTPD
fi

# 2. Sync blades to host. This is less straightforward, for several reasons:
# (i) We need root-SSH access to blades (password-free if you want decent
#     quality synchronization).
# (ii) We'd like to use nzipush ("nzipush -all sh date -u -s $BBDATE") but can't,
#     because nzipush doesn't project root privileges. Hence the escape to SSH.
# (iii) BusyBox doesn't recognize the normal date format (+%m%d%H%M%C%y.
#     %S),
#     so we have to RTFM to figure out what it does recognize.
function spuDates() {
    bbDATE=$1; shift
    while [ $# -gt 0 ]; do
        HWID="$1"; NAME="$2"; IPADDR="$3"; shift 3 # we only really need the
IPADDR here
        echo "Sync'ing BLADE=$HWID, NAME=$NAME, IPADDR=$IPADDR"
        ssh root@$IPADDR date -u -s $bbDATE &>/dev/null &
    done
    wait
}

# Do the work.
# TODO: DOES THE FOLLOWING MAGIC/UNDOCUMENTED COMMAND WORK
BETTER?
# nzipush -a -user root exec "date -u -s $(date -u +%C%y%m%d%H%M.%S)"
SPU_INFO=$(nzhw show -detail | grep '^SPU' | sed -e 's/^SPU *([0-9]*) *(\spa[0-9]*)\.\spu[0-9]*)' . * Ip Addr: \([0-9\.]*)\.*$/\1 \2 \3/)
[[ $($({#SPU_INFO[@]} % 3)) == 0 ]] || { echo '*** ERROR: NUMBER OF ELE-
MENTS IN SPU_INFO NOT MULTIPLE OF 3'; exit 42; }
BBDATE=$(date -u +%C%y%m%d%H%M.%S) # BusyBox special format
spuDates $BBDATE "${SPU_INFO[@]}"

# 3. Report.
# We'd like to use system time in sec.nano, but BusyBox date command doesn't
support +%s.%N.
DATE=$(date); DATESECNANO=$(date +%s -d "$DATE")
echo "Host : $DATE = $DATESECNANO"
nzipush -all sh date # BusyBox date command doesn't support +%s.%N

```

- ►My announcement to the effect that “Wahoo is 4x faster than Skimmer” (Appendix G.a) occurred here (first thing in the morning of June 8).◄

- From: Devesh Agrawal  
To: netezza-cambridge

Date: 06/08/2011 08:37 AM  
Subject: Working from M'bro today

In case any one is looking for me -- I will be in Marlborough today.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/08/2011 09:25 AM  
Subject: Re: Working from M'bro today

Are you using Wahoo2? Given the good news about TPC-DS, I'd like to re-run using turbo this time on Wahoo1, and also run Atomics on Wahoo2 to get Steve the info he needs.

■ From: Jerrold Richard Title/Cambridge/IBM  
To: netezza-cambridge  
Date: 06/08/2011 09:39 AM  
Subject: "update" is working in Wahoo

"update" is now working in Wahoo. For example:

```
> create table t (a int, b int);
> begin;
> insert into t values(1,1);
> insert into t values(2,2);
> insert into t values(3,3);
> commit;
> select * from t;
```

A	B
1	1
2	2
3	3

```
> update t set b=9999 where a=2;
UPDATE 1
> select * from t;
```

A	B
1	1
2	9999
3	3

I did some testing with somewhat larger and more complicated examples than the above, and I have tested on real hardware as well as VSPU. Seems to work.

"delete" was already working previously. It turned out that the Postgres code al-



ready implemented "update" as an insert followed by a delete, so basically all I needed to do was transplant my ndeddelete.cpp changes into ndeupdate.cpp and it worked.

I have checked in these changes.

One cautionary note: The current implementation of Wahoo "delete" and "update" needs to be optimized in its use of temp table space. I'm going to do that next. I'm not sure the current implementation scales up to really large deletes or updates. Maybe hold off on switching the perf tests to use update-based ballooning, until I do those optimizations.

■ From: Walter Tuvell  
To: Jerrold Richard Title  
Cc: netezza-cambridge  
Bcc: Daniel Feldman  
Subject: Re: "update" is working in Wahoo

I agree about postponing updating to the latest. Now that we've got a usable build, we want to collect usable numbers. And when I do update, I'll be sure to retain the one I'm using now.

But above and beyond all that, THANKS for getting DELETE/UPDATE working in Wahoo!

■ From: Devesh Agrawal  
To: Walter Tuvell  
Date: 06/08/2011 12:04 PM  
Subject: List of settings

Hi Walt, Can you send me your list of registry settings and I can edit it and send it over to Paul ?.

PS: Can I confirm that genstats was enabled for this tpcds run you did last night ?

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/08/2011 12:19 PM  
Subject: Re: List of settings

The settings were listed in the "Notes" section of the text file (both of them) that accompanied the graphic this morning, but I've also attached a more detailed version here.

And yes, genstats was enabled.

►Attached here was file “whereWeStand.txt”, with the following contents:◄

Base:

```
401_create_tpcds_db -s 100 --noballoon --force --debug -r n
runBar.py --test=tpcds --tpcdsdb=tpcds100 --force --debug -r n
```

Options:

```
--parallel=1 -- NO (i.e., do parallel)
genAllDbStats -- ON
updDsid -- OFF
redistrib_web.sql -- OFF
--skipstaging -- ON (that is, no staging)
JIT-stats -- OFF
Hammer-mode -- OFF
WLM/GRA -- OFF
2-phase-planner -- OFF
stats-collection -- OFF
```

Where:

- 0) "--parallel=1" is an option to the 401 script command-line.
- 1) "genAllDbStats" is a function in lib/nz\_functions.py. To turn it OFF, insert a return statement at the beginning.
- 2) "updDsid" is a function in the 401 script. To turn it OFF, insert a return statement at the beginning.
- 3) "redistrib\_web.sql" occurs in 1 place in the 401 script. To turn it OFF, comment-out that line and its 2 surrounding lines (and add a "pass" to keep python happy).
- 4) "--skipstaging" is an option to the 401 script.
- 5) JIT stats can be turned off by a global (postgres) option setting:  
     /nz/data/postgresql.conf:  
         enable\_jit\_stats=0  
         nzsystest restart (or: pkill -HUP -f postmaster)  
     or on a local (PerfBar) per-request basis:  
         Bar5.1/cfg/nzsql\_query.cfg:  
         set enable\_jit\_stats=0;
- 6) Hammer-mode is something from WLM/GRA-land that may be buggy. The DBOS waits for SPU to finish, the SPU does finish, but DBOS never gets notified.

To turn it OFF:

```
nzsqa showregistry
nzsqa set
host.schedFlags=16
```

- 7) WLM/GRA: Somehow perfbar user is getting only 1% resources?!

To turn it OFF:

```
nzsqa showregistry
nzsqa set
host.schedAllowGKandGRA=no # default=no
```

- ```
host.schedGRAEnabled=no    # default=yes; no means use GK
                           # instead of GRA
```
- 8) 2-phase planner can be turned OFF like JIT-stats, on a local (PerfBar) per-request basis:  
Bar5.1/cfg/nzsql\_query.cfg:  
set enable\_2phase\_planner=0;
- 9) Stats-collection is an option on runBar.py. To turn it off:  
./runBar.py --stats=n ...

■ From: Walter Tuvell  
To: Steve Lubars  
Date: 06/08/2011 12:45 PM  
Subject: Formatting

Steve, which do you like?

►Attachments omitted here; irrelevant to this Complaint. Their content consisted of formatting/protocol that Steve and I were working out for special "Steve" file to be included in "WaltBar" tests.◄

■ From: Walter Tuvell  
To: Steve Lubars,  
Date: 06/08/2011 01:44 PM  
Subject: Check these, please.

►Attachments omitted here; irrelevant to this Complaint. Their content consisted of formatting/protocol that Steve and I were working out for special "Steve" file to be included in "WaltBar" tests.◄

■ From: Devesh Agrawal  
To: Steve Lubars, Walter Tuvell  
Cc: Fritz Knabe  
Date: 06/08/2011 10:59 PM  
Subject: vt-disklog

Seems to work for me atleast when I do a simple scan of a 100G store-sales table, I get as many IO's in there as I expect. We can try it with your 'except' test query tomorrow.

May I confirm one thing : In the except test, there are only two snippets doing disk-io. So when walt says that he only sees 47 sec in the vt-disklog -- could it be the case that both of those snippets take up only about 47 seconds ?. IIRC, that may well be the case.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Cc: Fritz Knabe, Steve Lubars  
Bcc: Daniel Feldman  
Date: 06/09/2011 05:58 AM  
Subject: Re: vt-disklog

Devesh, what are you talking about here? Why do you think there's a problem with vtdisklog? I know nothing about this kind of a problem with vtdisklog, and nobody has mentioned anything about it to me, not even you, so why am I CC'd on this note?

What you say is perfectly correct. The vtdisklog captures Violin activity during a test, which in the case of something like EXCEPT is much shorter than overall test time, the remainder of the time being taken up by SPU. This was pointed out by me explicitly, several times to several people (for sure including you and Fritz, maybe not Steve because he wasn't very close to this kind of thing at that time), when I first distributed my ~20 perf reports on Atomics, over a month ago ("PerfTest-May06").

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/09/2011 08:27 AM  
Subject: More evidence

It occurred to me, belatedly, that maybe I should distributed the original (unre-touched) PerfBar summary reports for this morning's runs, so that any of you can check the results too. I invite that. After all, I was as excited as anybody that Wahoo looked fast, which no doubt contributed to my goof, so a cool third-party observer might help.

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/09/2011 08:29 AM  
Subject: More evidence

Damn, failed to attach the files (thanks to that fine HCI of Lotus Notes).

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Devesh Agrawal  
To: Fritz Knabe, Walter Tuvell, Steve Lubars  
Date: 06/09/2011 09:15 AM

Subject: vt-disklog sizing.

So here is how vt-disklog currently works at a high level: There is a userspace process that has mmap'd a nzlocal fixed size file. This file is the 'buffer' b/w the kernel and the 'select \* from vt-disklog' command. Userspace fills up this buffer by draining the kernel buffers as new data appears. And when we read the -vt-disk-log, this file is read and dumped to the virtual-table's recordstore.

Since this file is mmap'd, there is a cap on how many vt-disklog entries we can capture. This cap is configurable via a registry setting. Currently a vt-disklog entry occupies about 20 bytes and the current limit is 8M. So this means about 160MB of 'virtual memory' could be used. We can bump this limit up. But I find that bumping it up to say a 1GB or so usually is not possible since the address space is tight on our spu processes anyway. So the vt-disklog should currently show you only 8M entries but you can change that by the spuDisklogEntries cfg paramter.

Andy had suggested me a way to revamp this userspace vt-disklog code to not use mmap and just use regular file io (since this is sequential file access, regular file buffering would do as well as mmap). It will take me about 2 days to revamp this and that would remove this artificial limit on the disk-log size (barring nzlocal disk space). If you would like I can give this a shot.

►This email expressed Devesh's suggestion on how to "capture more vt-disklog". It seems to be correct, but it's suboptimal, because it's only a one-off solution for vt-disklog, requiring too much work in the wrong place. Subsequently, I solved this problem correctly when I wrote the nzVtCapture.sh script, because that solves the problem for all virtual tables, not just the vt-disklog virtual table. I informed Devesh of nzVtCapture.sh at that time. Appendix O.◄

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/09/2011 10:07 AM  
Subject: Screenshot of IM

I'm sending it in email because you're not online.

►Transcription of contents of screenshot of IM transcription follows here:◄

(09:15:49 AM) Devesh Agrawal: yesterday, right?

(10:02:44 AM) Walt: See email about calling perfmeasure.

(10:02:45 AM) Unable to send message: User is not online

(10:04:14 AM) Walt: What exactly do you want to know about the tpcds tests? I.e., what does "read only" mean? Once I know, I can figure them out.

(10:04:14 AM) Unable to send message: User is not online

(10:05:53 AM) Walt: Wahoo build is from Jun 7, see info in the .png and .txt. I'm also using out build system's convention: If kit name contains "dbg" it's debug, otherwise turbo.

(10:05:54 AM) Unable to send message: User is not online

(10:06:20 AM) Walt: (Rich told me not to update yesterday, because of the stuff he's been submitting.)

(10:06:20 AM) Unable to send message: User is not online

■ From: Walter Tuvell  
To: Daniel Feldman  
Date: 06/09/2011 11:38 AM  
Subject: The perfmeasure bug

Remember that perfmeasure bug of Devesh I mentioned to you at our 1-on-1 this morning? Well, it's not exactly a "bug", he just happened to change the calling sequence and forgot to tell me. So I've fixed it, and will soon try it out (once the current batch of tests finishes).

FWIW, this kind of stuff happens all the time, developers working things out amongst themselves, it's not a big deal, not something that should be surfaced to mgmt. So Fritz doesn't know about it, I see no reason to tell him, except possibly to "make sure I'm teflon" (which until now I didn't think I needed to do), and the only reason you know about it is this note.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/09/2011 12:20 PM  
Subject: The answer to your question

Xterm capture:

►Attachment omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/09/2011 02:40 PM  
Subject: qlist

A little better, but only very little. We're now 1.614x slower than Skimmer, instead of 1.676x.

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/09/2011 03:05 PM  
Subject: Turbo vs. debug builds

Q: How much faster is a turbo build over a debug build?

A: About 10%.

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/09/2011 04:275 PM  
Subject: Does Wahoo run as fast as itself?

I wanted to know if 2 tests of Wahoo (turbo) ran at the same speed. These are the 2 turbo runs, one from last night, one from this morning.

Judge for yourself.

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: Paul Houlihan  
Date: 06/09/2011 07:46 PM  
Subject: Wahoo settings

Paul, I know Devesh gave you a set of settings I'd put together for Wahoo (because he asked me for it), but I now notice I gave him an outdated copy. So here's the latest (mostly topy fixes). Sorry!

►File "whereWeStand.txt" (63 lines) was included here; it's already included above in this Appendix E (see the email dated 06/08/2011 12:19 PM).◄

■ From: Devesh Agrawal  
To: netezza-cambridge  
Date: 06/09/2011 08:17 PM  
Subject: Some more (low-priority) experiments to consider to understand why we



regressed.

I was thinking about what could be causing the slowdown from our earlier 0.9x ETR over skimmer vs the current 1.6x since March 31st. We have mainly done the following things since then: (a) Async Writing (b) 2 huge merges (c) Piecemeal GC stuff. (d) Updates, deletes (e) Lots of fpm related refactoring.

I posit that none of which should have any bearing on performance of the main players: CBS, Emu and gencode except the 2 merges. So I am wondering if we can try out these two experiments: (i) int-trunk on skimmer (ii) nonwahoo build of dev-wahoo on wahoobox. And compare both of these to the skimmer 6.0 basis that Walt has been using so far. In particular I am interested in whether the simple read only queries (which don't create any temp-tables or mvviews) perform over these.

Non-wahoo build of dev-wahoo on wahoobox would also be interesting because it will share our hardware and emu. We may also learn something useful by randomly comparing the plans produced by wahoo and skimmer. And finally, we can retry our good old scan performance experiments to see if we have introduced any bottlenecks in cbs and emu.

Just my 2 cents ...

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/10/2011 10:57 AM  
Subject: "WaltBar" TPC-DS stats

Folks, I bring your attention to wahoo2:/nz/PerfTest=wahoo2=2011-06-09=TPC-DS@100. It contains a run from last night of "WaltBar" output. In it are the files Steve uses for his FFM ("Fudge Factor Model", see the files names "steve\*"), as well as the standard perfReport=wahoo2=2011-06-09=QUERYxxx\*.txt files of the form that I've distributed in the past (on May 6, to be specific, that last time we had a working environment).

In the May report, I only distributed Atomics results (and only ~20 of them at that), not TPC-DS. But additionally, I have now also collected new Atomics results. They're at wahoo1:/nz/PerfTest=wahoo1=2011-06-09=ATOMICS@100.

These are all exploded views, not tarballs. I plan to leave them there indefinitely, for reference.

Note that there are some tests missing, because they failed for some reason yet-to-be-analyzed. Indeed, there may even be some failures in the bunch I haven't detected yet (I haven't yet perused them all).

Finally, note that spuperflog info is missing, because the script that gathers that info somehow didn't work. But that's probably of interest to only very few folks,

and we can re-run the tests when/if those folks need/want that info.

Enjoy!

- ►In the following, there are a number of technical emails between myself and Devesh. Their content will seem somewhat disjointed and out-of-context. That's because much of our conversation late on the afternoon of Friday, June 10, was conducted via IM, for which I have no record. These communications are all of a technical nature, whose content irrelevant for this Complaint. The point is not their content, but their existence. (Nevertheless, they're presumably archived on some IBM document retention server, and are discoverable there.)◄

- From: Walter Tuvell  
To: Devesh Agrawal  
Bcc: Daniel Feldman  
Date: 06/10/2011 04:17 PM  
Subject: Wahoo archeology

OK, this is going to be a complicated, so bear with me. What I'm going to prove is that wahoo\*/wahoobox/nz/Bar5.1-Mar31 is "mostly" the thing you want, that is, the Wahoo build upon which the Mar 30 "It's Alive" report on the whiteboard is based. Specifically, query015-099 works for you.

Our tool to prove this is going to be PerfScore.

To begin with, start with the numbers used in the It's Alive report. These numbers are given in the attachment called perfBar=WahooProto:Mar30ItsAlive=TPC-D-S@100=2011-03-31.csv (and its corresponding .txt file, included for human readability). Please look at those, and convince yourself they really are what I claim. I.e., compare them with the document posted on the whiteboard. (Note these are not original PerfBar "summary" files, I don't have those, never thought I'd need them).

Next, consider the file wahoo\*/wahoobox/nz/Bar5.1-Mar30/results/kit.6.1.D2.wahoo10\_wtuvell-Mar30/kit.6.1.D2.wahoo10\_wtuvell-Mar30.tpcds100.summary. That file is also included here, but renamed as perfBar=WahooProto:Mar30Build=TPC-DS@100=2011-03-31.summary.csv. Please verify that too.

So now, to use PerfScore, we need to have the files in a common format (same metadata, i.e., column structure). For that purpose, edit the 2 .txt's to put them into the right shape, calling them .txtA's, and then transform those into .csvA's. (It's only necessary to retain the tpcds info, you can throw away atomics etc., but it doesn't hurt to leave other stuff in because perfScore.py will notice and do the right thing.)

You're with me so far, right? OK. Now issue the following command:

```

perfScore.py --numerNamVer=WahooProto:Mar30Build \
--denomNamVer=WahooProto:Mar30ItsAlive \
--benchNamSiz=TPC-DS@100 \
--dateOfTest=2011-06-10 \
--numerFil=perfBar=WahooProto:Mar30Build=TPC-DS@100=2011-03-31.summary.csvA \
--denomFil=perfBar=WahooProto:Mar30ItsAlive=TPC-DS@100=2011-03-31.csvA \
--benchColNam=1:tpcds \
--testNamCol=2 \
--eTimeCol=3 \
--graphWidth= \
--graphHeight= \
--graphFont= \
--graphPointSize= \
--iDelim= \
--oDelim= \
--notes="Wahoo Archeology"

```

After you do that, you end up with the 2 perfScore=\*.png|txt] files.

Open the .png in an image viewer, and viola! You can see that everything from query015 to query099 is identical!

Therefore, anything you find in that range in the specified Bar5.1-Mar31 tree on wahoobox is indeed what you're looking for.

So what happened with query001-014? Easy. Look at the timestamps in perfBar=WahooProto:Mar30Build=TPC-DS@100=2011-03-31.summary.csv. It shows that those test results came from a test written after the one we're hunting for (that's the way PerfBar works, it overwrites entries in the summary file). I don't remember now what that test was, but from the looks of things it was obviously slower than the build we're hunting for. And since it stops after query014, it seems it was a buggy build (since I don't think I would voluntarily stop a test run in mid stream).

How's that for sleuthing?

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: netezza-cambridge  
Bcc: Daniel Feldman, Sujatha Mizar  
Subject: 1000 failed

So, creating the TPC-DS@1000 database failed. See wahoo1:/nz/Bar5.1/initialize/scripts/nohup-Jun10.out.

■ From: Devesh Agrawal  
To: Walter Tuvell  
Date: 06/10/2011 03:33 PM  
Subject: runBar.cfg

Felix suggests us to leave the --stats=y by default and instead hack around the individual stats options in runBar.cfg. That way we can pick and choose what we want. Look at this chat ...

dagrawal@us.i... Hey Felix. What's an easy way to turn off stats collection (all those virtual tables that we save in perfbar). But retain saving the plans ?.

<- saving the plan files I mean. 3:14:45 PM

Felix Santiago you can turn all those things off in runBar.cfg located in the cfg directory 3:16:43 PM

dagrawal@us.i... I find that turning off the stats also turns off the plan file saving. 3:17:30 PM

Felix Santiago is there a particular one that also turns off the plan collection? 3:18:20 PM

dagrawal@us.i... Hmm. So is this runBar.cfg a more finer grained control than simply running Bar with --stats=n. Perfbar is not a easy script to understand . 3:24:58 PM

Felix Santiago that is correct 3:25:19 PM

<- once you read through the code a couple of times, it starts making some sense, but I agree, it's not the easiest thing to wrap your head around 3:25:55 PM

dagrawal@us.i... So you recommend leaving --stats=y and then hacking the vars in the runBar.py ?. Is there a description of what each configuration does. I only want to avoid saving certain problematic virtual tables for wahoo. 3:26:05 PM

Felix Santiago yes, leave --stats=y (it's y by default BTW) and just turn off the things in runBar.cfg 3:26:51 PM

dagrawal@us.i... thanks ! 3:26:56 PM

Felix Santiago if you want to more details on some specifics, feel free to ask 3:27:28 PM

<- i'll try to answer what i can 3:27:35 PM

dagrawal@us.i... thanks ! 3:27:39 PM

Felix Santiago FYI, all the entries that begin with dump\_ just do a select \* from

the table 3:28:01 PM

<- everything that starts with perfqry\_ actually runs sql located in sql/perf\_capture  
3:28:26 PM

dagrawal@us.i...thanks. this is useful to know 3:28:46 PM

■ From: Walter Tuvell  
To: Devesh Agrawal  
Bcc: Daniel Feldman, Sujatha Mizar  
Date: 06/10/2011 04:32 PM  
Subject: Re: runBar.cfg

So, maybe I don't understand why you want to do this, that is what does PerfBar bring to the party? Most of the planfiles you're interested are already captured by "WaltBar", and are in wahoo2:PerfTest=wahoo2=2011-06-09=TPC-DS@100. The place to look is in the perfReport\*.txt's there.

■ From: Walter Tuvell  
To: Devesh Agrawal  
Bcc: Daniel Feldman, Sujatha Mizar  
Date: 06/10/2011 04:35 PM  
Subject: Re: runBar.cfg

Oh sorry, I was too fast on the trigger. I meant to explain what I meant by "most". Some of the tests failed to produce reasonable perfReport files, I'm not sure, that will be investigated. But most of them are OK, so you can look at those first, and if you find a smoking-gun trail we can then get the rest one way or the other (PerfBar or WaltBar)

■ From: Walter Tuvell  
To: Devesh Agrawal  
Bcc: Daniel Feldman, Sujatha Mizar  
Date: 06/10/2011 05:26 PM  
Subject: Mar 31 Skimmer

Devesh, you've asked for info about the Mar 31 Skimmer PerfBar results (the ones used in the "It's Alive" comparison with Wahoo), so they can be used to compare planfiles.

Well, I can't give that to you. I've searched and can't find them (didn't think at the time I'd ever need them, there was no reason to think Wahoo would take such a dive).

However, I can give you the next-best thing. By dumb luck, there's a PerfBar tree

at S1-8:nz/Bar5.1-Mar16 that never got cleaned up. By using the same sort of forensic techniques I sent you earlier for Wahoo, I have determined that it's morally equivalent to (though not exactly equal to) the It's Alive Skimmer run. The proof is given in the PerfScore output attached below.

So you can use that.

With this, I believe I have satisfied the 3-request IM you sent me earlier in the day. There is one additional (separate) request you sent me in a separate IM, and I will work on that next.

►Attachments omitted here; irrelevant to this Complaint.◄

■ From: Walter Tuvell  
To: netezza-cambridge  
Bcc: Daniel Feldman, Sujatha Mizar  
Date: 06/10/2011 06:00 PM  
Subject: The slowdown

Devesh asked if I'd publish a PerfScore report of the Wahoo slowdown, i.e, WahooNow (Jun 7) vs. WahooItsAlive (Mar 31). It's attached.

This, together with the failure of Wahoo to produce a 1000G TPC-DS DB, and some other work I've been doing this afternoon on-the-side with Devesh, leads me to believe that "almost" all the steps of the "upcoming performance tests" that Fritz proposed a couple days ago (but distributed to only a subset of us, so you may not have seen it).

By "almost", I mean that the host stats (nzmon) isn't yet done, and I'm not sure of Steve's status with his Fudge-Factor-Model. Of these, Steve's part seems to be doable without difficulty, if judging from his processing to date is a reliable guide. And, the host stats is very unlikely to yield any fruit, because I've been visually monitoring that for some time now (via those nice little gkrellm towers), and I've never yet seen anything in the host CPU or network graphs that seems at all suspicious.

And so with that, I sign off.

Cheers, everyone!

►Attachments omitted here; irrelevant to this Complaint.◄

## F Email Chain: Next Steps (June 2)

■ From: Fritz Knabe  
To: Steve Lubars  
Cc: Walter Tuvell  
Date: 06/02/2011 02:11 PM  
Subject: Next steps

Although we're waiting for progress on the Wahoo development build to be able to run TPC-DS, we should still be able to do some other investigation that was second in priority. With Walt's nzmon and other tools, we should now examine what's happening on the host when we do an atomics run and see those mysterious gaps where the SPU is effectively idle. There is also the question of whether the sawtooth in Violin I/O is real or an artifact of measurement. Steve, would you be able to start making some progress on this? Walt, I know you're mostly out of the office through the end of the week, but Steve may need a couple pointers to stuff to be able to make progress.

■ From: Steve Lubars  
To: Fritz Knabe  
Cc: Walter Tuvell  
Date: 06/02/2011 02:30 PM  
Subject: Re: Next steps

Are you suggesting a run of Atomics on a non-Wahoo build on a non-Wahoo machine? That should give insight into the SPU idle periods, though not necessarily the sawtooth pattern.

Walt, is your Skimmer (S1-8) available? Presumably we could get good data from another platform as well (Dwarf, ThinkTank, virtual ThinkTank).

■ From: Fritz Knabe  
To: Steve Lubars  
Cc: Walter Tuvell  
Date: 06/02/2011 03:48 PM  
Subject: Re: Next steps

I thought we could run atomics on a Wahoo build on a Wahoo machine, as Walt was able to run these successfully recently - if he still has that build. That is, whatever issue that is killing TPC-DS doesn't seem to be manifesting in the atomics run.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Steve Lubars



Date: 06/02/2011 03:51 PM  
Subject: Re: Next steps

Yes, I was able to run Atomics. Which WahooN can I use to see if it still works?

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Steve Lubars  
Date: 06/02/2011 03:52 PM  
Subject: Re: Next steps

Hold on, let me check the WahooN's to see which one I need (as I don't have the kit on my machine anymore, I transferred it).

■ From: Steve Lubars  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 06/02/2011 03:55 PM  
Subject: Re: Next steps

Hi Walt - they're both free at the moment.

■ From: Walter Tuvell  
To: Steve Lubars  
Cc: Fritz Knabe  
Date: 06/02/2011 04:02 PM  
Subject: Re: Next steps

It looks like I only have the May 24 build on those machines (and on my laptop, which is my build machine), which is known-bad for TPC-DS. And it's been so long, I don't remember to what extent it's good, do you? Did it boot? Did it create a TPC-DS 100G DB? I'll start exploring these questions now (on Wahoo1), and get back to you.

■ From: Steve Lubars  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 06/02/2011 04:11 PM  
Subject: Re: Next steps

I thought that 5/25 took a long time to generate 100GB. Although we know it's bad for TPC-DS, I didn't realize we knew it was good for Atomics.

Which build was used for the Atomics that generated the spreadsheets (and charts) I've been working with?

■ From: Fritz Knabe  
To: Steve Lubars  
Cc: Walter Tuvell  
Date: 06/02/2011 04:14 PM  
Subject: Re: Next steps

I think those results were dated May 6, so it must be an older build than that.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Steve Lubars  
Date: 06/02/2011 04:21 PM  
Subject: Re: Next steps

Right, May 6, and I don't have that around anymore.

I just tried booting NPS on Wahoo1, didn't work, will try the heavier hammer (re-boot with reinit) and see if that works.

■ From: Steve Lubars  
To: Walter Tuvell  
Cc: Fritz Knabe  
Date: 06/02/2011 04:54 PM  
Subject: Re: Next steps

Rich might have an April 27th build that is just what we want - it doesn't run TPC-DS, but should be fine for Atomics.

What's a little confusing is that Rich thinks that build contained a GC bug - but perhaps that wasn't tickled by Atomics? Or perhaps it invalidates some of our data....

■ From: Walter Tuvell  
To: Steve Lubars  
Cc: Fritz Knabe  
Date: 06/02/2011 04:59 PM  
Subject: Re: Next steps

I got my build to boot with reinit, will now try to run it. In the meantime, why don't you give me a pointer to Rich's build?

Talk to you tomorrow (I'll be at-home again, a little woozy after colonoscopy in morning).

■ From: Walter Tuvell  
To: Steve Lubars  
Cc: Fritz Knabe  
Bcc: Daniel Feldman  
Date: 06/03/2011 07:29 AM  
Subject: Re: Next steps

I started the test overnight (cron job), and this morning it was hung on the nested\_loops test, which is pretty basic (copied below). So I've now removed nested\_loops and restarted. Will report on that when I get back from endoscopy center, probably after noon.

```
nested_loops.sql:
SELECT COUNT(*)
FROM web_sales c1
INNER JOIN date_dim ON (c1.ws_sold_date_sk=d_date_sk)
INNER JOIN warehouse c2 ON (c1.ws_warehouse_sk>=c2.w_warehouse_sk)
WHERE d_year BETWEEN 2000 AND 2001;
```


## G Report And Email Chain: Wahoo 4× Faster Than Skimmer (June 8-9)

►The PerfScore report of June 8 consists of a page of graphics and a page of numerics, both reproduced here. The size of the original graphic (as posted on the whiteboard in Cambridge, for example) was a full 8½"×11" landscape, but is here reduced in size for documentation purposes.

Later (the night of June 8), this report was discovered to be in error, and the email chain reporting that error is included here as well (without its supporting PerfScore report, because that's irrelevant for the purposes of this Complaint).◄

## G.a PerfScore Report (June 8)

PERFScore — Wahoo1:Jun7DgbBuild / Skimmer:6.0.P7  
 TPC-DS@100 (2011-06-07)



### PERFScore REPORT

=====

Numerator Name:Ver = Wahoo1:Jun7DgbBuild  
 Denominator Name:Ver = Skimmer:6.0.P7  
 BenchMark Name@Size = TPC-DS@100  
 Date-Of-Test = 2011-06-07

#### Notes:

This now gives us our first REAL comparison between Wahoo and Skimmer!  
 But lots of asterisks: Debug build; GenStats ON; updDsid OFF;  
 redistrib\_web.sql OFF; JIT-stats OFF; Hammer-mode OFF; WLM/GRA OFF;  
 2-phase-planner OFF; stats-collection OFF.

| TestNum | TestNam   | NumerET | DenomET | ETR   | Rank | TestNum | TestNam  | NumerET | DenomET | ETR   | Rank |
|---------|-----------|---------|---------|-------|------|---------|----------|---------|---------|-------|------|
| ..... 1 | query001  | 17.650  | 71.260  | 0.248 | 62   | .... 52 | query050 | 25.650  | 235.750 | 0.109 | 15   |
| ..... 2 | query001v | 12.240  | 2.220   | 5.514 | 94   | .... 53 | query051 | 132.020 | 888.710 | 0.149 | 36   |



|         |            |          |           |       |      |         |            |          |           |       |      |
|---------|------------|----------|-----------|-------|------|---------|------------|----------|-----------|-------|------|
| .... 3  | . query002 | 423.160  | 1117.190  | 0.379 | . 79 | .... 54 | query051v  | 101.370  | 902.550   | 0.112 | . 18 |
| .... 4  | query002v  | 391.090  | 626.540   | 0.624 | . 83 | .... 55 | . query052 | . 19.360 | . 139.510 | 0.139 | . 31 |
| .... 5  | . query003 | . 14.240 | . 141.390 | 0.101 | . 11 | .... 56 | . query053 | . 27.510 | . 141.540 | 0.194 | . 54 |
| .... 6  | . query006 | . 35.010 | . 633.940 | 0.055 | . 3  | .... 57 | . query054 | . 41.530 | . 294.800 | 0.141 | . 33 |
| .... 7  | . query007 | . 50.920 | . 144.590 | 0.352 | . 75 | .... 58 | . query055 | . 20.070 | . 140.580 | 0.143 | . 34 |
| .... 8  | . query008 | . 19.670 | . 144.840 | 0.136 | . 29 | .... 59 | . query056 | . 71.390 | . 292.690 | 0.244 | . 61 |
| .... 9  | . query011 | 426.120  | 2206.100  | 0.193 | . 52 | .... 60 | query056v  | 65.810   | 286.400   | 0.230 | . 59 |
| .... 10 | query011v  | 199.610  | 1212.620  | 0.165 | . 43 | .... 61 | . query057 | . 37.110 | . 114.580 | 0.324 | . 73 |
| .... 11 | . query012 | . 17.040 | . 55.190  | 0.309 | . 71 | .... 62 | query057v  | 60.470   | 311.210   | 0.194 | . 53 |
| .... 12 | . query013 | 240.620  | 160.530   | 1.499 | . 88 | .... 63 | . query058 | . 63.850 | . 210.380 | 0.303 | . 69 |
| .... 13 | query014b  | .....    | .....     | ..... | .... | .... 64 | query058v  | 52.930   | 205.610   | 0.257 | . 66 |
| .... 14 | query014bv | .....    | .....     | ..... | .... | .... 65 | . query060 | . 54.930 | . 294.320 | 0.187 | . 51 |
| .... 15 | . query015 | 263.700  | 206.180   | 1.279 | . 85 | .... 66 | query060v  | 44.510   | 289.720   | 0.154 | . 39 |
| .... 16 | . query016 | . 19.430 | . 222.310 | 0.087 | . 7  | .... 67 | . query061 | . 41.960 | . 285.270 | 0.147 | . 35 |
| .... 17 | . query017 | .....    | .....     | ..... | .... | .... 68 | . query062 | . 18.980 | . 51.380  | 0.369 | . 76 |
| .... 18 | . query018 | . 42.380 | . 112.590 | 0.376 | . 78 | .... 69 | . query063 | . 25.480 | . 140.390 | 0.181 | . 50 |
| .... 19 | . query019 | . 28.790 | . 171.520 | 0.168 | . 44 | .... 70 | . query065 | 143.230  | 1399.140  | 0.102 | . 12 |
| .... 20 | . query020 | . 16.990 | . 95.870  | 0.177 | . 46 | .... 71 | . query066 | . 32.780 | . 150.850 | 0.217 | . 57 |
| .... 21 | . query021 | . 37.110 | . 22.850  | 1.624 | . 89 | .... 72 | query067b  | 31.010   | 359.080   | 0.086 | . 5  |
| .... 22 | . query022 | 421.840  | 255.530   | 1.651 | . 91 | .... 73 | . query068 | . 28.810 | . 160.370 | 0.180 | . 47 |
| .... 23 | query022v  | 415.210  | 252.950   | 1.641 | . 90 | .... 74 | . query069 | . 40.190 | . 294.950 | 0.136 | . 30 |
| .... 24 | . query025 | .....    | .....     | ..... | .... | .... 75 | . query070 | . 90.980 | . 281.280 | 0.323 | . 72 |
| .... 25 | . query026 | . 28.720 | . 98.020  | 0.293 | . 68 | .... 76 | . query071 | . 31.850 | . 288.350 | 0.110 | . 17 |
| .... 26 | . query027 | . 19.480 | . 145.710 | 0.134 | . 28 | .... 77 | . query072 | .....    | .....     | ..... | .... |
| .... 27 | . query028 | . 33.000 | . 834.230 | 0.040 | . 1  | .... 78 | query072v  | .....    | .....     | ..... | .... |
| .... 28 | . query029 | .....    | .....     | ..... | .... | .... 79 | . query073 | . 18.590 | . 141.850 | 0.131 | . 26 |
| .... 29 | . query030 | . 25.930 | . 172.770 | 0.150 | . 37 | .... 80 | . query074 | . 75.200 | . 686.990 | 0.109 | . 16 |
| .... 30 | query030v  | 26.570   | 282.970   | 0.094 | . 9  | .... 81 | query074v  | 73.460   | 624.480   | 0.118 | . 20 |
| .... 31 | . query031 | . 87.010 | . 231.510 | 0.376 | . 77 | .... 82 | . query076 | . 46.990 | . 289.150 | 0.163 | . 42 |
| .... 32 | query031v  | 81.430   | 641.150   | 0.127 | . 25 | .... 83 | . query077 | . 49.990 | . 312.430 | 0.160 | . 41 |
| .... 33 | . query032 | . 22.980 | . 192.480 | 0.119 | . 22 | .... 84 | query077v  | 40.180   | 303.590   | 0.132 | . 27 |
| .... 34 | . query033 | . 75.440 | . 295.170 | 0.256 | . 64 | .... 85 | . query079 | . 24.200 | . 157.830 | 0.153 | . 38 |
| .... 35 | query033v  | 70.200   | 289.970   | 0.242 | . 60 | .... 86 | . query081 | . 25.290 | . 150.540 | 0.168 | . 45 |
| .... 36 | . query034 | . 16.890 | . 142.300 | 0.119 | . 21 | .... 87 | query081v  | 28.330   | 240.920   | 0.118 | . 19 |
| .... 37 | . query036 | . 36.280 | . 143.240 | 0.253 | . 63 | .... 88 | . query082 | . 14.760 | . 160.400 | 0.092 | . 8  |
| .... 38 | . query037 | . 13.920 | . 114.900 | 0.121 | . 23 | .... 89 | . query083 | . 42.760 | . 56.250  | 0.760 | . 84 |
| .... 39 | . query038 | . 55.520 | . 517.030 | 0.107 | . 13 | .... 90 | query083v  | 32.540   | 52.610    | 0.619 | . 82 |
| .... 40 | . query039 | 788.770  | 340.460   | 2.317 | . 92 | .... 91 | . query084 | . 38.580 | . 28.700  | 1.344 | . 87 |
| .... 41 | query039v  | 158.380  | 118.980   | 1.331 | . 86 | .... 92 | . query086 | . 25.760 | . 54.130  | 0.476 | . 80 |
| .... 42 | . query040 | . 35.350 | . 667.450 | 0.053 | . 2  | .... 93 | . query087 | . 69.090 | . 563.130 | 0.123 | . 24 |
| .... 43 | . query041 | . 13.860 | . 52.670  | 0.263 | . 67 | .... 94 | . query089 | . 31.920 | . 144.050 | 0.222 | . 58 |
| .... 44 | . query042 | . 19.850 | . 140.970 | 0.141 | . 32 | .... 95 | . query090 | . 19.450 | . 98.410  | 0.198 | . 56 |
| .... 45 | . query043 | . 27.870 | . 141.350 | 0.197 | . 55 | .... 96 | . query091 | 100.320  | . 37.640  | 2.665 | . 93 |
| .... 46 | . query044 | . 31.930 | . 559.500 | 0.057 | . 4  | .... 97 | . query092 | . 45.490 | . 293.230 | 0.155 | . 40 |
| .... 47 | . query046 | . 28.660 | . 159.400 | 0.180 | . 48 | .... 98 | . query094 | . 16.480 | . 164.520 | 0.100 | . 10 |
| .... 48 | . query047 | . 60.650 | . 179.510 | 0.338 | . 74 | .... 99 | . query096 | . 15.030 | . 139.800 | 0.108 | . 14 |
| .... 49 | query047v  | 128.280  | 500.130   | 0.256 | . 65 | ... 100 | . query098 | . 25.730 | . 141.990 | 0.181 | . 49 |
| .... 50 | . query048 | . 81.250 | . 147.350 | 0.551 | . 81 | ... 101 | query099   | 29.340   | 95.900    | 0.306 | . 70 |
| .... 51 | . query049 | . 29.090 | . 336.660 | 0.086 | . 6  | .. ~Sum | .....      | 7283.390 | 28400.040 | ..... | .... |

ETR\_Stat| Value  
.. Length|... 101  
.. Count|.... 94  
... gMin| . 0.040  
.... gQ1| . 0.124  
... gMed| . 0.180  
.. \*gMean| . 0.224  
.... gQ3| . 0.308  
... gMax| . 5.514  
... gRan|139.380  
... gIQR| . 2.489  
... gMAD| . 1.521  
.... gAD| . 1.973  
... gStd| . 2.509

## G.b Follow-Up Email Chain (June 8-10)

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/08/2011 09:58 PM  
Subject: New numbers, and a correction

I'm unhappy to report that the "4x faster" numbers reported earlier today were in error. I accidentally compared Wahoo 100G against Skimmer 1000G (a one-digit fat-finger).

My sincere apologies. (Let him who is bug-free throw the first stone.)

The corrected number is "1.6x slower". See attached. And yes, I'll run it all yet again to double-check.

►Attachments omitted here; they're already included as the two pages of graphics and numerics immediately above.◄

■ From: Walter Tuvell  
To: netezza-cambridge  
Cc: Daniel Feldman  
Date: 06/09/2011 07:49 AM  
Subject: Double-check of Wahoo v. Skimmer

As promised, I re-ran the TPC-DS@100 perf test, the results are attached. They confirm last night's finding: Wahoo is 1.676x slower than Skimmer.

To ensure correctness, I generated fresh runs on the 2 machines concurrently, and monitored the results, by grepping the PerfBar output log for "Start test" at the beginning of lines. These are also attached. As you can see by looking at the time-stamps, I actually gave Wahoo a headstart by starting it first, but nevertheless Skimmer caught up quickly and surpassed it.

I was very careful, and I strongly believe there is no mistake this time.


►Attachments omitted here; irrelevant to this Complaint.◄

## H Email Chain: Upcoming Performance Tests (June 8-10)

■ From: Fritz Knabe  
To: Devesh Agrawal, Walter Tuvell, Steve Lubars  
Date: 06/08/2011 05:09 PM  
Subject: Upcoming performance tests

Here is the current prioritized list. Comments welcome! There may be a better order to explore these. I especially would like to know if any of these tests are burdensome to run because something hasn't been built out yet; in that case, we can make decisions about what to revise so that we're not wasting too much time on manual steps.

Should be  
bullet list, see  
Add. III, p. 5-6



Use turbo build to run TPC-DS on 100 GB with monitoring of SPU resources and no NPS ("perfbar") stats collection. Even though the host resources will not be measured, this will allow us to get the early picture of how Wahoo is comparing against TwinFin, normalized on a rack-to-rack basis. This run will not provide much information that is useful for performance problem diagnosis, but it will give us a high level view. Also, it will show us whether the SPU resource monitoring is slowing down the test by any appreciable amount, as we will be able to "mostly" compare this test with the one Walt ran on the night of June 7, where no monitoring was turned on. I note "mostly" because the June 7 run was on a debug build, not a turbo build. Other notes:

We will do this run without disk\_log stats. Previously, Walt found that generating these did not impact the performance much, but we may want to examine how the system behaves now, plus we may not have the right machinery to gather all the stats. So we will not have this data, and that we will therefore have no cache simulation. This will make our performance look worse than it should.

Although we want monitoring of host resources, we are only not measuring them in this test because the tools are still under development.

If this test can be run the night of June 8 and analyzed the morning of June 9 with Steve's scripts, I will use the information in the status meeting with Arvind and Netezza managers in the afternoon of June 9.

Use turbo build to run TPC-DS on 100 GB with monitoring of SPU resources turned on, host resource monitoring turned on if available, NPS/perfbar stats, and the disk log. This will give us all the info necessary for diagnosis and (perhaps) for running the cache simulator. However, we will probably run this test before host monitoring tools are available, so not all diagnosis will be possible.

The disk log machinery is in question right now. My understanding is that the disk log is automatically truncated when it reaches a certain size, and it may be that the 100 GB TPC-DS tests exceed that size. However, even if the disk log is just a sub-



set, that can still be used as a sample that Steve can extrapolate from. We will want to build the right disk log machinery, whether it is a patch by Devesh that expands the max size of the log, or a tool that dumps it periodically, but that is not a precursor to this test.

Ideally this test can run simultaneously with the previous test, on the other Wahoo VM.

Use turbo build to run atomics on 100 GB with monitoring of SPU resources (plus host is available) and no other stats collection. This will provide a quick picture on how well we're doing on atomics. If we have previously found that resource monitoring is affecting the test times, we will want to do this run both with and without monitoring. That will become yet another scaling factor that we have to work in. I'm hoping that this won't be necessary, or that we can brainstorm ways to change our monitoring so that it imposes less overhead.

We will rerun this test with host monitoring turned on as soon as that capability is present.

Use debug build to run atomics on 100 GB with all measurement, monitoring, and disk log. This will allow performance diagnosis as well as running the cache simulator to correct the results from the turbo run.

Repeat all of the above, with 1 TB data set. I anticipate we'll hit some problems, perhaps during building, and we may already have other tests we want to run before we get to this.

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Devesh Agrawal, Steve Lubars  
Bcc: Daniel Feldman  
Date: 06/10/2011 12:01 PM  
Subject: Re: Upcoming performance tests

If I understand this list correctly, (& Steve) have now completed "almost all" of items 1-3.

The remaining parts of 1-3, as well as 4, need some clarification. There are some suggestions of a debug build, but I don't understand what that would accomplish. The only reason I was doing that was to get through the 2.5 weeks of showstopping bugs. There's also a suggestion of running PerfBar with --stats=y, but Devesh has told me there's no need for that, because WaltBar already contains all the stats we require (modulo the spuperflog bug).

So I'll await clarification on these issues, and in the meantime try building the 1000G DB we need.

■ From: Walter Tuvell  
To: netezza-cambridge  
Bcc: Daniel Feldman, Sujatha Mizar  
Date: 06/10/2011 6:00 PM  
Subject: The slowdown

Devesh asked if I'd publish a PerfScore report of the Wahoo slowdown, i.e, WahooNow (Jun 7) vs. WahooItsAlive (Mar 31). It's attached.

This, together with the failure of Wahoo to produce a 1000G TPC-DS DB, and some other work I've been doing this afternoon on-the-side with Devesh, leads me to believe that "almost" all the steps of the "upcoming performance tests" that Fritz proposed a couple days ago (but distributed to only a subset of us, so you may not have seen it).

By "almost", I mean that the host stats (nzmon) isn't yet done, and I'm not sure of Steve's status with his Fudge-Factor-Model. Of these, Steve's part seems to be doable without difficulty, if judging from his processing to date is a reliable guide. And, the host stats is very unlikely to yield any fruit, because I've been visually monitoring that for some time now (via those nice little gkrellm towers), and I've never yet seen anything in the host CPU or network graphs that seems at all suspicious.

And so with that, I sign off.

Cheers, everyone!

►Attachments (2-page PerfScore report) omitted here; irrelevant to this Complaint.◄

## I Email Chain: Yelling Incident (June 8-10)

### I.a Comments On Some “Miscommunications” (June 8)

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Daniel Feldman  
Date: 06/08/2011 09:42 PM  
Subject: Comments on some "miscommunications"

Fritz, for this "iteration", my main task as been to generate TPC-DS PerfScore comparisons of Wahoo against Skimmer, for DB sizes 100G & 1000G.

As you know, I've tried many times (approx 20?) to run the 100G test over the past

2.5 weeks, and it's been impossible because of bad software. During that time, I have stated several times, both in person (e.g., stand-ups) and in email, that I was running debug builds of Wahoo first (because it's come in handy several times for debugging), and then once that succeeded I'd run a turbo build. Everyone agreed to this. The debug build finally succeeded last night, and I reported it this morning (erroneously, see separate email). At the time I spoke to you this morning when I posted the results, I reiterated that it was a debug build, with no stats generated (these things were even stated explicitly in the PerfScore report), and that I'd run the turbo build tonight. This was also repeated a couple of hours later, at today's stand-up. But at about 3pm, you jumped on me for not having run a perf-stat-ready build (i.e., turbo), with stats, so that Steve could use the stats. You did so publically, in the Camb office, in a LOUD voice, in a CONDESCENDING manner (you know, like the time you publically berated Michael Sporer when he and everyone else became impatient when you fumbled for 15 minutes at the Wahoo status meeting with a recalcitrant presentation [surely you must have sensed everyone looking at one another about the inappropriateness of the outburst, even John Metzger commented on it aloud after the meeting]). And you further stated explicitly that you thought I had gathered the stats for Steve, even though I had communicated the contrary to you clearly several times as stated above. And you even said aloud that verbal communication wasn't working between you and I, so you'd start writing things down, even though it is provably (as just proved) you who are doing the communicating! This was not acceptable behavior from my point of view, and I asked you to get off my back.

Another thing you said about my effort over the last 2.5 weeks was that "since you haven't finished the host CPU gathering work yet (i.e., nzmon), it seems the bad Wahoo software isn't slowing you down". That was an absurd statement. For, 80% of what I've been doing during that time was trying to run the many Wahoo builds (for that was higher priority than getting nzmon working and integrated), with only 20% devoted to nzmon. It is very, very obvious that Wahoo itself has been the long pole, not me.

You also tasked Steve for this iteration to come up with a FFM ("Fudge Factor Model"), from data generated by me. For 2 weeks, Steve labored in essential isolation from me on that task, coming up with the model, and figuring out how to implement it. Finally this week he asked me about parts of what he was doing, and I explained to him that I was already doing it all, with the exception of applying his FFM model itself (i.e., I was already gathering the stats and boiling it down to exactly what was needed for input to the FFM). He was blown away. In other words, whatever it was you communicated to him (almost all of which was in my absence), it did not include any of the work I'd already done. Yesterday, I heard him tell you that "Walt is already doing 95% of what I need", and today I heard him tell you "Walt is already doing 97.5% of what I need". In other words, there was a clear miscommunication between you and Steve, but once he started talking to me we communicated well enough to get him 97.5% of the way towards his goal. Why didn't you get me involved with Steve earlier, it would have saved him a lot of time? For example, do you remember when Steve announced at stand-up on Monday that he was trying to get nmon running on the host, and then immediately afterwards when it was my turn I announced I already had nzmon running on the host? Does it seem there may be

a pattern of miscommunication here?

For this iteration, my secondary task is to feed Steve some numbers, as mentioned above. As is well-known, the only such numbers I've gathered in the past has been ~20 tests from Atomics, and it's been stated several times that was where Steve was to start, then finishing Atomics, then moving on to TPC-DS. As has been stated consistently at stand-ups, Steve was getting basically nowhere until this week, and then when he and I started talking he became much more upbeat, and we both said we were working out the "protocol" (i.e., content and formatting issues for me to send him the info so he could feed it to his FFM). And indeed as late at TODAY'S stand-up, we both stated that we were still working out this protocol, and when you pressed the issue about when you'd see some actual results, I said that we'd send you "at least one sample" by end-of-day today. I in fact did send Steve the info he needed to do that at 1:30 today, telling him that he could use it to generate the FFM output you wanted. (Whether or not he did send that to you I don't know.) Yet, at the aforementioned 3pm outburst today, you stated that you had expected to see a complete set of FFM results! You were wrong to state these things at all, much less in a public forum, because you had all the correct information readily at your attention yet you didn't heed it.

## I.b Follow-Ups And Apologies (June 10)

■ From: Walter Tuvell  
To: Fritz Knabe  
Cc: Daniel Feldman  
Date: 06/10/2011 10:16 AM  
Subject: Re: Comments on some "miscommunications"

Just one final (important) fact that needs to be put on the record about this incident at this time (while it's fresh, lest future recollections be in error).

Fritz, at that 3pm conversation on Wed, when I reminded you of the debug/no-stats mode of the Wed morning test (stated to you already twice earlier that day), you said (after a moment's hesitation) that you thought I meant no PerfBar-stats had been gathered, but that I had gathered my "WaltBar"-stats (the ones that get sent to Steve). It was at that point you raised your voice even above its normal volume, and I responded by asking you to get off my back.

Unfortunately I was unable to summon up a more appropriate reaction at that moment, because my automatic, reactive fight-or-flight instinct kicked in. For, what you said was so nonsensical that I couldn't even logically parse it in the instant given to me. Namely as you know full well (because we've talked about it multiple times), the PerfBar-stats are collected by the PerfBar-harness, while the WaltBar-stats are collected (and auto-edited) by the WaltBar-harness (which is Wahoo-specific, because it includes Devesh's PerfMeasure and SpuPerfLog). So it makes NO SENSE WHATSOEVER to think that the WaltBar-stats could be generated by a PerfBar run (whether no-stats or otherwise), much less massaged into the shape

Steve needs. The PerfScore graphic and numerics I had posted that morning are generated from PerfBar runs (not from WaltBar runs), of course. This is universally well-known to anyone who has read the PerfScore spec, which you have, because you've commented upon it to me, and you even corrected a misstatement of mine about the PerfScore graph that very morning, as were stating together at the white-board studying it.

That illogic was so off-the-wall that I was unable to analyze it in-the-moment (I only connected the dots Wed night, tossing in bed, after I sent the the first note in this chain), and it's why the only reaction I came up at 3pm with was to go into a defensive posture, asking you to halt your attack (as I perceived it). So, even though my voice stayed consistently lower than yours, I do hereby offer my apologies for raising my voice enough to make myself heard.

I also wish to offer you and your team the best possible success going forward (desire for that success is indeed what prevented me from "smelling a rat" when I posted the wonderful 4x-better results that morning), and I will of course give Dan & Sujatha my full cooperation and support for the transition.

I'm sorry it's come to this, because I truly thought I was extremely fortunate to be in the best possible project at Netezza. And I hope someday we can rehash all this at leisure and come to a better understanding of what's happened to us, and patch up the rift. I'm just sad that day cannot be today.

■ From: Fritz Knabe  
To: Walter Tuvell  
Cc: Daniel Feldman  
Date: 06/10/2011 12:09 PM  
Subject: Re: Comments on some "miscommunications"

Walt, I do appreciate the words of encouragement for Wahoo as we roll forward and your support for the ongoing performance work. I am of course disappointed with how our working relationship on this project has come apart, and I'm sorry for the embarrassment and anger I provoked when raising my voice. Clearly we were both frustrated! However, I'm hopeful that you will be able to have a much better experience on your next project.

## J Email Chain: Request For Help (June 10)

■ From: Walter Tuvell  
To: Kelli-ann McCabe  
Date: 06/10/2011 10:27 AM  
Subject: A request for help

Kelli, you and I haven't met yet (I joined Netezza last Nov), but Dan Feldman point-

ed me to you as someone I may be able to talk to.

I am forwarding to you the transcript of an interaction that's happened to me this week. I would like to learn how I could have handled things better.

Please feel free to contact Dan and/or Fritz if you wish.

►Attachments omitted here; they're already included in this Complaint (first two emails in Appendix I).◄

■ From: Kelli-ann McCabe  
To: Walter Tuvell  
Date: 06/10/2011 11:15 AM  
Subject: Re: A request for help

Hi Walter,

I am happy to assist and welcome a discussion. I will not be in this afternoon, but could arrange to talk on Monday afternoon? If not, I could talk Wednesday, Thursday or Friday next week.

■ From: Walter Tuvell  
To: Kelli-ann McCabe  
Date: 06/10/2011 11:19 AM  
Subject: Re: A request for help

Thank you Kelli-ann (sorry I mispronounced your name), Mon afternoon would be fine, feel free to schedule a Notes mtg for anytime if you like.

I really appreciate this.

■ From: Walter Tuvell  
To: Kelli-ann McCabe  
Date: 06/12/2011 9:46 AM  
Subject: Fw: Weekly report

Kelli-ann, it occurs to me that I should inform you about what I wrote to Dan in my weekly report.

See you tomorrow.

►Attachment omitted here; it's already included in this Complaint (Appendix A.dd).◄

■ From: Walter Tuvell  
To: Kelli-ann McCabe  
Date: 06/13/2011 10:43 AM  
Subject: Explanation, to some extent

Kelli-ann I've talked to my wife (Linda) about the matter I'm involved in, of course, and she suggested (wisely) that I offer you a few words of explanation. I'm not trying to change the meaning of anything I've written, just to explain what the words themselves and presentation mean.

1. The weekly report I sent you was initially intended to be "just" a weekly report, not a "piece of evidence" (that is, I hadn't thought to CC you initially). For that reason, it was written in a manner consistent with the usual communications I have with Dan. Since he's always kept up-to-date with my work throughout the week anyway, the weekly report is merely a quick-and-dirty summary, written in a sort-of "shorthand" so as to avoid wasting anyone's time. Just as we all sometimes use colorful language with one another in our one-on-one meetings (to no one's offence or objection), I do the same in my weekly reports. We can dredge up all my weekly reports from email records if necessary to check this point, but I suspect I am not an exceptional case about intended-private conversations such as this. So: my apologies if anything came across an unnecessarily candid, it was nothing Dan hadn't already heard before, and this was just a summary. It was certainly not "hot-headed" or "abrasive" in any way, when seen in this light of a one-to-one conversation between participants with a history of communicating amicably this way. (And before this morning's surprising response by Dan, I am certain he would have said the same thing.) It was merely "boys being boys" (if I may use that admittedly sexist but commonly-used colorful phrase).

2. In particular, if the weekly report seemed to be overly adversarial, it was not intended that way. For, up until now, I had always thought Dan and I had an unusually clear-channel of communication. We've had occasional minor misunderstanding of course, but those have always been "mere wording" issues, pounced upon immediately and resolved to a common agreement as to meaning. Again, I could come up with several illuminating examples (can't we all?). Indeed this recent claim of Dan's that I've misrepresented what he's said may be another instance of the same thing, and I have to admit I was shocked by his bald statement without details so that I could offer a clarification if needed.

3. Particular attention may be paid to the word "worthless" in the weekly report, that was something I'd already said to Dan, and he understood I meant in the strict meaning of "without substantive value". However, it is not a word I've used outside private discussion with Dan. Instead, I will forward to you immediately following this note the original 2 emails I sent about this topic, and you will see it nowhere uses the word "worthless" or anything else inflammatory. One of these emails is entitled "Re: Upcoming performance tests" (it includes Fritz's original note outlining the "upcoming performance tests" he wanted done), and the other is entitled "The slowdown". These might have some contextual reference to other "mysterious" things, but if you point those out to me I can clarify them to you (and/or forward further emails about them as appropriate).



4. Referring to "adverse job action" and "Siberia" were certainly not intended to "dis" any other project going on in this company! I think I made that clear in my earlier note this morning, but it's worth repeating, that what I was referring to was the "killing somebody in the middle of the night" as being an extremely pejorative action to anyone's career (as must certainly be common knowledge). I do admit that I viewed the Wahoo project as the best possible position for me. In fact it was Dan himself who told me it was the plum position for anybody to nab (the way he phrased it was "there's almost no other job like this for a performance professional in the country"). But to say Wahoo was the "best" is not to say other projects are "bad".

5. The place I used the phrase "minor niceties", well yes, that was bit smarmy. But as I did also already point out in my note earlier this morning, I'd mentioned this "lack of acknowledgement of justice to Walt" to Dan previously (in our meeting of ~3 weeks ago), and he refused at that time to acknowledge the sentiment, instead repeating his mantra of "only interested in the success of Wahoo". For reference, when I talked to John Metzger a few minutes ago, he spoke in terms of "primary responsibility to success of all his Advanced Development projects, but also with a major interest in the welfare of all his people, and a fun place to work". I'd expected something like that from Dan too, but was sadly disappointed.

See you this afternoon.

►The mentioned "original two emails" forwarded to Kelli-ann omitted here, they're already included elsewhere in this Complaint (they're the emails of Appendix I; I ended up sending all three of them, not just the first two).◄

## K Sujatha's Work Items

Sujatha Mizar  
6/10/2011

### Project Status

---

#### 1. Pushing Runtime Restricts to the FPGA:

- Gathered Hash Join statistics and sent the data to Garth and Dan for review. Need to send it to a larger audience (John Yates, Bill Ackerman, Jason Viehland and Michael Sporer) after incorporating comments from Garth and Dan. Raw data is on `/mount/snap1a/smizar/HashJoinStatsData`

I can send you a note that has instructions on how to create the necessary DB and tables and load the data.

- Working on coming up with the Performance Criteria for sub-select and in-list queries. Once this is done need to update the Performance Criteria document.

#### 2. Twinfin midlife kicker:

- Loaded Oprofile data from a PerfBar run into NZ80533. Need to start analyzing this data.
- Need to re-run Perfbar with Oprofile set up for gencode assembly annotation.

#### 3. Dev-Emu instrumentation:

- Currently running a TPCDS1000B run on NZ80748-H1. Will be post-processing the data into a DB for Garth.

#### 4. Block IO tracing:

- Ran a Proof-Of-Concept test with a single query for Dan and collected the blktrace/blkparse data. The query output and the blktrace/blkparse output is in: `/mount/snap1a/smizar/test_blktrace`
- Need to run blktrace/blkparse with a multi-user Perfbar test suite and post-process the data into a DB for analysis.

## L Email Chain: Hail And Farewell (June 11-13)

►*The goodbye letter, entitled "Hail and Farewell", was emailed to my six Cambridge colleagues. The content was the same in each case, and was addressed to them individually. Later everything (both the letter and the responses) was also sent to Kelli-ann McCabe (06/14/2011 12:55 PM).*◄

### L.a Goodbye Letter

■ To: ►*Devesh Agrawal, Dan Dietterich, Rich Title, Jeff Keller, Huamin Chen, Steve Lubars*◄

From: Walter Tuvell

Date: 06/11/2011 9:50 AM

Subject: Hail and Farewell

Dear ►*Colleague*◄, as you've no doubt heard, I'm leaving the Wahoo project, by "mutual agreement".

But to you personally I want to offer my apologies. It seems I have been reported to be a liar and a bully. At least that's what I've been told indirectly, for my accuser reported me without saying anything directly to me it at all. This doesn't feel good to me. I've of course asked for a 3-way (or more) conversation to clear the air, multiple times, but that's been adamantly refused. So I have absolutely no idea who I've offended, or how. It is for that reason I am forced to write this note in ignorance. If you're not one of the people I've offended, please ignore this note. But in case you're among the offended, please know I did not intend to do whatever it is I'm being accused of, nor was I even aware of it. Nevertheless I do profoundly apologize to you, quite sincerely. The irony is, I've always valued my interactions with you, as clear, honest and of high quality in all ways, and I've never detected any friction between us. If I've offended you, I really hope we can iron out our difficulties at some future date.

In any case, I truly do wish you all the best, both personally and professionally, and I certainly do hope Wahoo is successful in all ways. For what it's worth, I'm including below a transcript of an event that happened on Wednesday afternoon, which seems to have precipitated this turn of events. I've tried to be honest and accurate in what I wrote, but it's only my point of view. As mentioned above nobody else has told me what their point of view is (several people were in the office at the time). In any case, the transcript at least shows that Fritz and I are departing on the best terms possible under the circumstances, so I hope you can be forgiving as well.

►*Included here was the message stream contained in Appendix I.*◄

## L.b Responses

■ From: Devesh Agrawal  
To: Walter Tuvell  
Date: 06/11/2011 10:03 AM  
Subject: Re: Hail and farewell

Hi Walt,

I am not sure what gave you the idea that you offended me. Not in the slightest :). I have learnt some high quality stuff from you : (a) How to measure performance correctly and (b) How to write good scripts.

I was in Marlboro on wednesday so I was not party to the discussion you had with Fritz/Steve etc. All I know is that there was some problem with the vt-disklog stuff which I investigated and couldn't reproduce. And the last I heard about this whole saga was Fritz telling us yesterday that you won't be working on this project.

But anyway, let bygones be bygones. I hope you can bring some rigour and 'science' into the overall Perf group.

I will catch you later when I come to Mbro during our weekly wednesday pilgrimages. Take care

■ From: Walter Tuvell  
To: Devesh Agrawal  
Date: 06/11/2011 10:19 AM  
Subject: Re: Hail and farewell

Devesh, thank you for getting back to me so quickly.

I absolutely I do NOT think it was you who I offended, or complained about me! But as I pointed out, they won't tell me anything, so I felt the need to apologize to you, "just in case".

I, too, have very greatly enjoyed our working relationship, and I hope it can continue in some way! I, too, have learned a lot from you (And, no, the vtdisklog thing was not a part of any of this that I am aware of, it just happened to occur around the same time and I assume it was just all of us being innocently confused to some degree, except for you!) I envy (in a good sense) your knowledge of NPS inner workings, which I aspire to (but that's hard not being in core development), and the way you "follow" your nose to find the source of difficulties, like with the current perf problems.

Being new to the company, one of the things I've been struggling with is "figuring out what everybody sort-of knows already". That is, I'm not sure if I'd be silly to try

writing a new PerfScore tool, or if nobody cared about it. So I looked around, and didn't find anything in the company that satisfied my curiosity, so I decided I need to write it. Your recent interest in PerfScore is really the finest memory I have during this unpleasantness.

Anyway, I will indeed be located in Marl from now, and I certainly do wish to see you when possible. Remember that day I hugged you? I feel like doing that again now.

■ From: Daniel Dietterich  
To: Devesh Agrawal  
Date: 06/12/2011 8:00 PM  
Subject: Re: Hail and farewell

Hi Walt,

This is a bummer. I had not heard and I have not been offended.

I hope you find yourself on something fun and interesting.

Regards - and see you around, I hope!

■ From: Walter Tuvell  
To: Daniel Dietterich  
Date: 06/13/2011 7:56 AM  
Subject: Re: Hail and farewell

Thank you Dan, I of course did not for a second believe you were any part of my problem, but I had to touch all bases because they won't tell me what's going on.

I hope I'll be seeing you too!

■ From: Steve Lubars  
To: Walter Tuvell  
Date: 06/13/2011 12:22 AM  
Subject: Re: Hail and farewell

The subject of your departure didn't come up until around 5pm on Friday; on the day of the incident, I may have said something to the effect of, "that was uncomfortable", but other than that I have no knowledge of whatever else may have been reported by whom.

Thanks for your well wishes (both personal and project). And in a sense we still work together - it just won't be several feet apart.

■ From: Walter Tuvell  
To: Steve Lubars  
Date: 06/13/2011 7:58 AM  
Subject: Re: Hail and farewell

Thanks Steve, as I said I have no doubt you are completely innocent here, but I had to write you to be sure, since they won't tell me what's going on. I don't remember you saying "that was uncomfortable" or anything else, ►Because I was, literally, in shock/PTSD.◄ I only remember you staring straight out the window trying to be invisible. Which was totally appropriate. This is not your fight.

Best wished, and be seeing you. I hope!

■ From: Huamin Chen  
To: Walter Tuvell  
Date: 06/13/2011 07:44 AM  
Subject: Re: Hail and farewell

I am really shocked and very sorry for your decision. I am not quite involved with the storage work and so haven't worked with you for most part of good past months. Still I have always regarded you as great colleague. If you feel uncomfortable with the experimental results, I don't think they hurt me. Nobody is exempt from mistakes like this and what is true for me is that your persistent pursuit of perfection matters most. I never regard you as any sort of "liar or bully", if you ever thought I would accuse people in such way.

I hope all the best for you and your work. Please do enjoy your work. We will meet each other still see each other in Marlborough some time

■ From: Walter Tuvell  
To: Huamin Chen  
Date: 06/13/2011 8:01 AM  
Subject: Re: Hail and farewell

Thank you for the kind words, Huamin. Please know I did not ever believe you and I had ever had crossed word with one another! However, they won't tell me what's going on, so I cannot know for sure.

It was not my decision to leave my decision to leave Wahoo, instead I was "fired" from Wahoo for reasons I do not understand (they won't tell me the reasons!). The incident of my mistaken report last Wed is not an issue, it was a normal human error, and I corrected it as soon as I discovered it. (But that was good guess!)

I really do wish you and Wahoo all the best, and I too hope to be seeing you in Marlboro.

■ From: Rich Title  
To: Walter Tuvell  
Date: 06/12/2011 5:54 PM  
Subject: Re: Hail and farewell

It certainly wasn't me saying anything like that. I've always found you to be pleasant to work with.

I do wish you the best!

■ From: Walter Tuvell  
To: Rich Title  
Date: 06/13/2011 8:11 AM  
Subject: Re: Hail and farewell

Rich, I never for a moment thought you and I had a problem, but they won't tell me what's going on so I had to touch all bases.

Of course I wish you and Wahoo all the best, and I'll be seeing you in Marlboro, I hope!

■ From: Jeffrey Keller  
To: Walter Tuvell  
Date: 06/13/2011 10:14 AM  
Subject: Re: Hail and farewell

Actually, I hadn't heard; this is a complete surprise to me. I'm sorry to hear it, and I assure you you've done nothing to offend me. Let me know how the new stuff goes; I hope it works out well.

■ From: Walter Tuvell  
To: Rich Title  
Date: 06/13/2011 10:48 AM  
Subject: Re: Hail and farewell

Jeff, please rest assured I have no thought whatsoever that there has been any friction between us! But since they won't tell me what's going on, I had to touch all bases.

Thank you for your kind sentiments, and look forward greatly to seeing you the next time you're in Marlboro. I hope!