

Saturday, May 22. 2010

Writing an OAuth Provider Service

Last year I showed how to use pecl/oauth to write a Twitter OAuth Consumer. But what about writing the other end of that? What if you need to provide OAuth access to an API for your site? How do you do it?

Luckily John Jawed and Tjerk have put quite a bit of work into pecl/oauth lately and we now have full provider support in the extension. It's not documented yet at php.net/oauth, but there are some examples in svn. My particular project was to hook an OAuth provider service into a large existing Kohana-based codebase. After a couple of iterations this should now be trivial for others to do with the current pecl/oauth extension.

Step 1 - Create a Consumer Key registration page

In order for an application to communicate with your service you assign it what is essentially a user id and password. Except in OAuth terms these are known as the Consumer Key (CK) and shared secret. Your main job here is to make sure the CK and secret are unique and unguessable which means you need a decent entropy source. On Linux you have `/dev/random`, and the non-blocking potentially slightly weaker `/dev/urandom` mechanism. You can check how much entropy is available with:

```
cat /proc/sys/kernel/random/entropy_avail
```

In general something like the following should give you a decent random string of characters that you can use for your CK and secret:

```
function new_consumer_key() {
    $fp = fopen('/dev/urandom','rb');
    $entropy = fread($fp, 32);
    fclose($fp);
    // in case /dev/urandom is reusing entropy from its pool, let's add a bit more entropy
    $entropy .= uniqid(mt_rand(), true);
    $hash = sha1($entropy); // sha1 gives us a 40-byte hash
    // The first 30 bytes should be plenty for the consumer_key
    // We use the last 10 for the shared secret
    return array(substr($hash,0,30),substr($hash,30,10));
}
```

You can of course read more entropy and use a longer hash, like sha256 or whirlpool if you want longer keys.

Step 2 - The OAuth endpoints

It would probably be a good idea to skim the OAuth Spec at this point. We need two main end points for the OAuth sequence of requests. The first is called the "request token" end point. In my case I named it `/v1/oauth/request_token`. And since I am sliding this into a Kohana-based system, I wrote a controller (`application/controllers/v1.php`) which contained this in the constructor:

```
try {
    $this->provider = new OAuthProvider();
    $this->provider->consumerHandler(array($this,'lookupConsumer'));
    $this->provider->timestampNonceHandler(array($this,'timestampNonceChecker'));
    $this->provider->tokenHandler(array($this,'tokenHandler'));
    $this->provider->setParam('kohana_uri', NULL); // Ignore the kohana_uri parameter
    $this->provider->setRequestTokenPath('/v1/oauth/request_token'); // No token needed for this end point
    $this->provider->checkOAuthRequest();
}
```

```
} catch (OAuthException $E) {  
echo OAuthProvider::reportProblem($E);  
$this->oauth_error = true;  
}
```

This makes the pecl/oauth extension do all the heavy lifting for us. We register a couple of callback functions. `lookupConsumer` will look up the CK and check if it is valid. `timestampNonceChecker` will check whether the timestamp of the request is sane and falls within the window of our Nonce checks. And this function will, of course, also check whether the provided Nonce has been used already to prevent replay attacks. And finally the `tokenHandler` callback will check whether a request or access token is valid. The `setParam('kohana_uri',NULL)` call tells the extension to ignore the `kohana_uri` parameter in the URLs because this was injected by an nginx rewrite rule and wasn't part of the original request.

The `lookupConsumer` callback happens early on. My version looks like this:

```
public function lookupConsumer($provider) {  
    $consumer = ORM::Factory("consumer", $provider->consumer_key);  
    if($provider->consumer_key != $consumer->consumer_key) {  
        return OAUTH_CONSUMER_KEY_UNKNOWN;  
    } else if($consumer->key_status != 0) { // 0 is active, 1 is throttled, 2 is blacklisted  
        return OAUTH_CONSUMER_KEY_REFUSED;  
    }  
    $provider->consumer_secret = $consumer->secret;  
    return OAUTH_OK;  
}
```

Just a simple lookup in the data model and if found, the shared secret is attached so it can be used to check the signature of the request.

So, now we need to write code for the different stages of the OAuth request. I have an `oauth` method with a switch:

```
public function oauth($action=NULL) {  
    if($this->oauth_error) return;  
  
    switch($action) {  
    case 'request_token':  
        $token = Token_Model::create($this->provider->consumer_key);  
        $token->save();  
        // Build response with the authorization URL users should be sent to  
        echo 'login_url=https://'.Kohana::config('config.site_domain').  
            '/session/authorize&oauth_token='.$token->tok.  
            '&oauth_token_secret='.$token->secret.  
            '&oauth_callback_confirmed=true';  
        break;
```

Remember, the CK and the signature of the request has been checked already in the constructor, so by the time we get here, as long as `$this->oauth_error` is false, we can generate a request token and respond with a urlencoded set of parameters that include the unauthorized request token and secret we just generated along with a login url that the 3rd-party application will redirect the user to in order for them to authorize the request token to be used on their behalf.

Step 3 - Authorizing the request token

Back to your regular web UI, you now need to add a landing page for authorizing request tokens. Make sure the language on the page makes it clear to the user that they are authorizing a 3rd-party application to act on their behalf. It is a good idea to make them re-enter their password and explicitly click a button to do this authorization.

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Step 4 - The Access token

Once the user has authorized the request token, the 3rd-party app needs to exchange the authorized request token for an access token. My switch case for this looks something like this:

```
case 'access_token':
    $access_token = Token_Model::create($this->provider->consumer_key, 1);
    $access_token->save();
    $this->token->state = 2; // The request token is marked as 'used'
    $this->token->save();
    // Now we need to find the user who authorized this request token
    $utoken = ORM::factory('utoken', $this->token->tok);
    if(!$utoken->loaded) {
        echo "oauth error - token rejected";
        break;
    }
    // And swap out the authorized request token for the access token
    $new_utoken = Utoken_Model::create(
        array('token'      => $access_token->tok,
            'user_id'     => $utoken->user_id,
            'application_id'=> $utoken->application_id,
            'access_type' => $utoken->access_type));
    $new_utoken->save();
    $utoken->delete();
    echo "oauth_token={$access_token->tok}&oauth_token_secret={$access_token->secret}";
    break;
```

As you can see I have a couple of data models that deal with tokens here. Token_Model is the token itself while UToken_Model links users to tokens. A token is either a request token (type 0) or an access token, (type 1) and they can be in various states. The tokenHandler callback keeps track of whether a token is used correctly. It looks like this:

```
public function tokenHandler($provider) {
    $this->token = ORM::Factory("token", $provider->token);
    if(!$this->token->loaded) {
        return OAUTH_TOKEN_REJECTED;
    } else if($this->token->type==1 && $this->token->state==1) {
        return OAUTH_TOKEN_REVOKED;
    } else if($this->token->type==0 && $this->token->state==2) {
        return OAUTH_TOKEN_USED;
    } else if($this->token->type==0 && $this->token->verifier != $provider->verifier) {
        return OAUTH_VERIFIER_INVALID;
    }
    $provider->token_secret = $this->token->secret;
    return OAUTH_OK;
}
```

Once we have handed out an access token, the 3rd-party app can store this and sign API requests with the associated shared secret and send that signature along with the consumer key. As you can see in the above token handler, there is a revocation check. At any point the user can go in through the Web UI and revoke an access token for a specific 3rd-party application.

Step 5 - An actual API call

Now that we have completed the OAuth dance, we can write API end points. Because all the oauth magic happens in our controller constructor, we are left to just implement our own api logic in our api methods. Mine look something like this:

```
public function user($cmd, $arg1=NULL, $arg2=NULL) {
    if(!$this->checkAccess()) return;
    $res = array();
    switch($cmd) {
    case 'settings':
        foreach($this->user->settings as $key=>$val ) {
            $res[$key] = $val;
        }
        $this->jsonResult($res);
        break;
    }
```

The checkAccess function uses the access token to look up the user who authorized it and sets \$this->user the same way your login flow will load the user once authenticated.

For the most part, the new OAuth Provider support in pecl/oauth takes the pain out of writing an OAuth Provider. It is smart about detecting the URL of the end points. You only have to indicate what the request_token (or 2-legged) endpoint is called, and it takes it from there. You can also pass the endpoint explicitly to checkOAuthRequest at each stage of the OAuth dance, but then you have to separate the steps more than I have here. I find this approach nice and compact and it lets me focus on the API part. I also like that it supports the OAuth ProblemReporting extension. This makes it a bit easier to track down problems. And by problems I mean signature mismatches. Everyone who has ever done anything with OAuth has struggled with them, and you will too. With the ProblemReporting extension you will get the raw string that the provider used to generate the signature. You just have to match that against what your consumer signed and you should be able to track down signature problems quickly.

Let me know if you do something interesting with the pecl/oauth extension. We may still be missing some use cases and right now there is decent momentum to fix things, so get in there and try it out.

Posted by Rasmus in PHP at 21:50

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Tuesday, February 9, 2010

A quick look at XHP

Facebook released a new PHP extension today that supports inlining XML. This is a feature known as XML Literals in Visual Basic. Go read their description here:
<http://www.facebook.com/notes/facebook-engineering/xhp-a-new-way-to-write-php/294003943919>

It adds an extra parsing step which maps inlined XML elements to PHP classes. These classes are core.php and html.php which covers all the main HTML elements. The syntax of those class definitions is a bit odd. That oddness is explained in the How It Works document.

Essentially, it lets you turn:

```
renderBaseAttrs() . ' />';  
}  
}
```

which extends html-element which in turn extends primitive. You can go read all the code for those yourself.

Note that to build XHP you will need flex 2.5.35 which most distros won't have installed by default. Grab the flex tarball and ./configure && make install it. Then you are ready to go.

I pointed Siege at my rather underpowered AS1410 SU2300 with the above trivial form examples. The plain PHP one and the XHP version. Ran each one 5 times benchmarking for 30s each time. The plain PHP one averaged around 1300 requests/sec. Here is a representative sample:

```
acer:~> siege -c 3 -b -t30s http://xhp.localhost/1.php  
** SIEGE 2.68  
** Preparing 3 concurrent users for battle.  
The server is now under siege...  
Lifting the server siege... done.  
Transactions:      38239 hits  
Availability:     100.00 %  
Elapsed time:     29.60 secs  
Data transferred:  3.97 MB  
Response time:    0.00 secs  
Transaction rate: 1291.86 trans/sec  
Throughput:       0.13 MB/sec  
Concurrency:      2.93  
Successful transactions: 38239  
Failed transactions: 0  
Longest transaction: 0.05  
Shortest transaction: 0.00
```

And the XHP version:

```
Transactions:      868 hits  
Availability:     100.00 %  
Elapsed time:     29.28 secs  
Data transferred:  0.08 MB  
Response time:    0.10 secs  
Transaction rate: 29.64 trans/sec
```

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Throughput: 0.00 MB/sec
Concurrency: 2.99
Successful transactions: 868
Failed transactions: 0
Longest transaction: 0.21
Shortest transaction: 0.05

So, a drop from 1300 to around 30 requests per second and latency from less than 10ms to 100ms. Running XHP on plain PHP is definitely out of the question. But, knowing that Facebook uses APC heavily and looking through the code (see the MINIT function in ext.cpp) we can see that it should play nicely with APC. So, re-running our PHP version of the form, now with APC enabled, that goes from 1300 to around 1460 requests per second, and no measurable latency:

Transactions: 43773 hits
Availability: 100.00 %
Elapsed time: 29.88 secs
Data transferred: 4.55 MB
Response time: 0.00 secs
Transaction rate: 1464.96 trans/sec
Throughput: 0.15 MB/sec
Concurrency: 2.93
Successful transactions: 43773
Failed transactions: 0
Longest transaction: 0.07
Shortest transaction: 0.00

The XHP version of the form now with APC enabled:

Transactions: 9707 hits
Availability: 100.00 %
Elapsed time: 29.45 secs
Data transferred: 0.94 MB
Response time: 0.01 secs
Transaction rate: 329.61 trans/sec
Throughput: 0.03 MB/sec
Concurrency: 2.97
Successful transactions: 9707
Failed transactions: 0
Longest transaction: 0.21
Shortest transaction: 0.00

Much better. But it is still around a 75% performance drop from 1460 to 330 and a ~10ms latency penalty. And yes, I did have a default filter enabled for these tests, so there was basic XSS filtering in place for the naked `$_POST['name']` variable in the plain PHP version. Of course, the default filtering would likely fail if the user data was used in a different context. And this 75% is obviously going to depend on what else is going on during the request. If you are spending most of your time calculating a fractal or waiting on MySQL, you may not notice XHP very much at all.

The bulk of the time is spent in all the tag to class interaction. If the core.php and html.php code was all baked into the XHP extension, it would be a lot quicker, of course. So, when you combine XHP with HipHop PHP you can start to imagine that the performance penalty would be a lot less than 75% and it becomes a viable approach. Of course, this also means that if you are unable to run HipHop you probably want to think a bit and run some tests before adopting this. If you are already doing some sort of external templating, XHP could very well be a faster approach.

Update: Here are the callgraphs.

The first is the plain PHP+APC version without XHP. And the second is the PHP+APC+XHP version. In the first you see all sorts of bits and pieces all across the stack getting cpu time.

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

In the second graph we see the effect of needing to copy and instantiate those core and html classes on every request. They are cached in APC, of course, but because of PHP's perfect sandbox, they cannot persist.

So we went from spending around 1% of our time in the executor to over 80%.

This isn't an entirely fair comparison, of course, since the plain version has close to no PHP to execute while the XHP version has 93 userspace classes to deal with. I would guess that XHP could get quite a boost if at least the primitives in core.php could be baked into the extension. Ideally all 93 basic html classes would be in C++ in the extension itself, but that would be a bit of a tedious undertaking.

Posted by Rasmus at 21:22

Thursday, February 4, 2010

HipHop PHP - Nifty Trick?

In a response to a question from ReadWriteWeb, among other things, I wrote:

My main worry here is that people think this is some kind of magic bullet that will solve their site performance problems. Generating C++ code from PHP code is a nifty trick and people seem to have gotten quite excited about it. I'd love to see those same people get excited about basic profiling and identifying the most costly areas of an application. Speeding up one of the faster parts of your system isn't going to give you anywhere near as much of a benefit as speeding up, or eliminating, one of the slower parts of your overall system.

The "nifty trick" part of that seems to have become the story, and them injecting a "just" in front of it makes it sound more derogatory. Anyone who knows me knows that I am a big fan of nifty tricks that solve the problem. When I first heard about the Facebook effort I was assuming they were writing a JIT based on LLVM V8 or something along those lines. Writing a good JIT is hard. Doing static code analysis and generating compilable C++ from it is indeed a nifty trick. It's not "just" a nifty trick, it is a cool trick that takes advantage of a number of characteristics of PHP. The main one being that you can't overload PHP functions. `strlen()` is always `strlen`, for example. In Python, this would be harder because you can overload everything.

I also noted that most sites on the Web have a lot of lower hanging fruit that would provide a much bigger performance improvement, if fixed, than doubling the speed of the PHP execution phase. The ReadWriteWeb site, for example, needs 160 separate HTTP requests and 41 distinct DNS lookups to load the front page. And once you get beyond the frontend inefficiencies you usually find Database issues, inefficient system call issues and general architecture problems that again aren't solved by speeding up PHP execution.

If you have done your homework and find that your web servers are cpu-bound, you are already using an opcode cache like APC and your Callgrind callgraph shows you that the PHP executor is a significant bottleneck, then HipHop PHP is definitely something you should be looking at.

Posted by Rasmus at 10:50

Sunday, January 10, 2010

SQLi Detection - Duh Moment

Not sure why it took me so long to figure out what I am sure is obvious to most other people who have thought about this, but it never clicked for me how to get anywhere near useful SQL Injection detection. The injection itself is trivial, of course, but determining whether it actually worked and weeding out false positives in an automated manner was something that seemed too hard.

During my run on Friday I had a Duh! moment on it. Annoyingly simple. Do it in 3 requests. Request #1 is a normal request. For example, "?id=1" in the URL. If the id is being passed to an SQL request it will return a single record or perhaps no record, it doesn't really matter. Now on request #2 do "?id=1 or 3=4", that is, inject a false 'OR' condition. If the output changes, we are done. Nothing to see here. However, if the output does not change we send request #3 with "?id=1 or 3=3" and if that output differs from request #2 then we have a potential SQLi situation. There are of course still chances of false positives (and negatives) with page stamps and such, but filtering out the response headers and html comments cuts down on that a bit. Add different combinations of single and double-quotes, like "?id=1'or'3='3" (without the double-quotes, of course) and it might be able to catch something.

The best thing about it is that it can slide into an existing scanner framework quite easily. If you have a base reference request, then it just adds a single request to the common case where the false 'OR' condition output does not match the base reference. You only need to do the true 'OR' condition request in case it does match.

Anybody have any other approaches?

Posted by Rasmus in Software at 18:44

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Thursday, September 24, 2009

Playing with Gearman

This was written in September 2009 when the current version of Gearman was 0.9. Thanks to Eric Day for answering my dumb questions along the way.

To get started, install Gearman. I am on Debian, so this is what I installed:

```
% apt-get install gearman gearman-job-server gearman-tools libgearman1 libgearman-dev libdrizzle-dev
```

Enable Gearman in /etc/default/gearman-server

Set up Gearman to use MySQL for its persistent queue store in /etc/default/gearman-job-server

```
PARAMS="-q libdrizzle --libdrizzle-host=127.0.0.1 --libdrizzle-user=gearman \  
--libdrizzle-password=your_pw --libdrizzle-db=gearman \  
--libdrizzle-table=gearman_queue --libdrizzle-mysql"
```

```
% mysqladmin create gearman
```

```
% mysql  
mysql> create USER gearman@localhost identified by 'your_pw';  
mysql> GRANT ALL on gearman.* to gearman@localhost;
```

** Careful, if you are running MySQL using --old-passwords this won't work with libdrizzle. You will need to get the 41-char password hash with a little snippet of PHP that does the double sha1 encoding:

```
% php -r "echo '*'.strtoupper(sha1(sha1('your_pw',true)));"
```

```
% mysql  
mysql> UPDATE mysql.user set Password='above_output' where User='gearman';
```

```
% mysqladmin flush-privileges
```

Then start the server.

```
% /etc/init.d/gearman-job-server start
```

Check to make sure gearmand is running. If it isn't, check for errors in /var/log/gearman-job-server/gearman.log

Also note that your password will be visible with 'ps' with this setup. The Gearman guys will be addressing this in gearmand-0.10

Next, let's get the gearman PHP extension installed. Grab it from svn:

```
% svn co http://svn.php.net/repository/pecl/gearman/trunk gearman  
% cd gearman  
% phpize  
% ./configure --with-php-config=/usr/local/bin/php-config
```

(If you have autoconf problems, apt-get install autoconf2.59 and set your

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

PHP_AUTOCONF env variable to "autoconf2.59")

% make install

edit your php.ini file and add: extension=gearman.so

Now we can add a worker. A worker is something that will process a Gearman request. We can write it in almost any language, but here is one in PHP (worker.php):

```
#!/usr/local/bin/php
```

Posted by Rasmus at 14:57

Monday, April 27, 2009

Using pecl/oauth to post to Twitter

I have seen a lot of questions about OAuth and specifically how to do OAuth from PHP. We have a new pecl oauth extension written by John Jawed which does a really good job simplifying OAuth.

I added Twitter support to Slowgeek.com the other day and it was extremely painless. The goal was to let users have a way to have Slowgeek send a tweet on their behalf when they have completed a Nike+ run. Here is a simplified description of what I did.

First, I needed to get the user to authorize Slowgeek to tweet on their behalf. This is done by asking Twitter for an access token and secret which will be stored on Slowgeek. This access token and secret will allow us to act on behalf of the user. This is made a bit easier by the fact that Twitter does not expire access tokens at this point, so I didn't need to worry about an access token refresh workflow.

First, a bit of DB groundwork that has nothing to do with OAuth itself. I needed a place to store the access tokens and secrets and relate them to the existing user ids. My Slowgeek user table has a numeric u_id field for each user, so that is what I am using as my primary key here:

```
CREATE TABLE twitter (  
  u_id int(10) not null,  
  name char(32) default NULL,  
  state smallint default 0,  
  token varchar(64) default NULL,  
  secret varchar(64) default NULL,  
  description varchar(255) default NULL,  
  status varchar(140) default NULL,  
  location varchar(80) default NULL,  
  followers smallint default 0,  
  mtime timestamp default CURRENT_TIMESTAMP on update CURRENT_TIMESTAMP,  
  PRIMARY KEY (u_id)  
) TYPE=MyISAM;
```

And the related DB code using PDO.

So, I now have a mechanism for storing Twitter-related data. Now for the real oauth work. First I registered my application with Twitter to get a consumer key and consumer secret. You do that at http://twitter.com/oauth_clients/new.

Because I am updating the status I needed read/write access for this app. Now the PHP code. It will probably be easier to read if you copy and paste it out of the iframe into your favourite editor.

The last thing this script did after it had gotten the access token and secret was to call `verify_credentials.json` to get the user's credentials.

That gives us json back which we can decode with `json_decode()`. There is some interesting stuff in your Twitter record. Here is mine from a few minutes ago:

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

stdClass Object

```
(
  [favourites_count] => 0
  [profile_text_color] => 666666
  [description] => Breaking the Web
  [screen_name] => rasmus
  [utc_offset] => -28800
  [profile_background_image_url] => http://static.twitter.com/images/themes/theme9/bg.gif
  [profile_link_color] => 2FC2EF
  [following] =>
  [profile_sidebar_fill_color] => 252429
  [url] =>
  [name] => Rasmus Lerdorf
  [time_zone] => Pacific Time (US & Canada)
  [protected] =>
  [status] => stdClass Object
  (
    [truncated] =>
    [in_reply_to_status_id] => 1642930101
    [text] => @DonMacAskill Floating point values are approximations in all computer languages #php #broken
    [in_reply_to_user_id] => 813491
    [favorited] =>
    [in_reply_to_screen_name] => DonMacAskill
    [id] => 1643107564
    [source] => Nambu
    [created_at] => Tue Apr 28 21:57:56 +0000 2009
  )

  [profile_sidebar_border_color] => 181A1E
  [notifications] =>
  [profile_background_tile] =>
  [followers_count] => 2112
  [friends_count] => 71
  [profile_background_color] => 1A1B1F
  [profile_image_url] => http://s3.amazonaws.com/twitter_production/profile_images/52489510/rl_normal.jpg
  [location] => Sunnyvale, California
  [id] => 928961
  [statuses_count] => 577
  [created_at] => Sun Mar 11 15:39:19 +0000 2007
)
```

The `$debug = $oauth->getLastResponseInfo();` after an oauth call will always give you info about the last call. For this verify_credentials call the LastResponseInfo is (with the actual token and secret deleted):

```
[u_id] => 1823955881
[name] => rasmus
[state] => 2
[token] => XXX
[secret] => XXX
[description] => Breaking the Web
[status] => @cdamian I don't like things strapped to my chest. My brain is a perfectly good built-in heart rate monitor.
[location] => Sunnyvale, California
[followers] => 2080
[utype] =>
[mtime] => 2009-04-26 09:07:08
```

And now you can actually send a tweet. This assumes the filename is 'twitter.php'. This shows how to POST a tweet to

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

twitter. Two things to note here. Because Twitter requires a status update to be sent as a POST request, we have to use OAUTH_AUTH_TYPE_FORM when we instantiate the oauth object here. And second, I have a CSRF-preventing crumb in the POST data. The idea here is to tie the POST body to the current user's login cookie so the bad guys can't spoof a form post to our twitter.php script.

I am hoping the code is for the most part self-explaining. You can also have a look at the great pecl/oauth examples.

Posted by Rasmus in PHP at 15:20

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Thursday, March 19, 2009

Select * from World

I have been having a lot of fun with two Yahoo! technologies that have been evolving quickly. YQL and GeoPlanet. The first, YQL, puts an SQL-like interface on top of all the data on the Internet. And the second, GeoPlanet, introduces the concept of a WOEID (Where-On-Earth ID) that you can think of as a foreign key for your geo-related SQL expressions.

First some example YQL queries to get you used to this concept of treating the Internet like a database. Go to the YQL Console and paste these queries into the console to follow along.

```
select * from geo.places where text="SJC"
```

This looks up "SJC" in GeoPlanet and returns an XML result containing this information:

```
12521722
Airport
Norman Y Mineta San Jose International Airport
United States
California
Santa Clara
```

```
Downtown San Jose
```

```
95110
```

```
37.364079
-121.920662
```

```
37.35495
```

Posted by Rasmus in PHP at 15:06

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Thursday, May 8, 2008

SearchMonkey

One of the things I have been playing with lately is Yahoo!'s SearchMonkey project. It appeals to me on many different levels. The geeky name is a play on GreaseMonkey. But instead of writing plugins that run locally in the browser, SearchMonkey is a way to write plugins for the Yahoo! Search results page that change the appearance of the results themselves. Best explained with an example. Assume I am looking for a Japanese restaurant, and on my search results page I see:

That's ok, I guess. It tells me it is somewhere in Redwood City and that it is a neighborhood restaurant, whatever that means. Compare that to:

This gets me a real address and phone number plus a number of other useful bits of information. That is the first level SearchMonkey appeals to me on. The usefulness is obvious. My usefulness test is to see if I can explain it to my mother. Having her search for recipes and get pictures of dishes, ingredients and preparation times right on the search results page makes this an easy sell.

The second level this appeals to me on is the way it is implemented. Writing these SearchMonkey plugins becomes much simpler if the site you are writing the plugin for uses microformats of some sort. hCard, hCalendar, hReview, hAtom, xfn or generic structured eRDF or RDFa tags. The data can also be collected via a separate XML feed that can then be converted via XSLT in the SearchMonkey developer tool. The microformat data is collected and indexed and when you go to write a plugin and specify the url pattern you are writing the plugin for, it will find whatever indexed metadata it has for that url. If it doesn't have what you are looking for, you can still write a custom data scraper to get it, but that gets a bit more involved. I really like that the easy path is to add some sort of semantic markup to the pages. Yes, as Micah points out, this is not the (uppercase) Semantic Web, but it is still a push towards semantic markup. Having such a tangible and visible result of adding semantic tags is going to encourage people other than microformat geeks to do so. The more semantic markup we get, the better off the Web is.

The third part that appeals to me is the way the plugins are written. You write a little snippet of PHP. It is actually a method in a class you can't see, but its job is to return an associative array of data such as the title to display, the summary, extra links to show and whatever other key/value pairs you might want in the output. Because you have a full-featured scripting language available, you can write quite complicated logic in one of these plugins and pull whatever data you want from the site the plugin is written for.

You can also write an add-on to your plugin which is called an Infobar. It is a little bar that is shown below the plugin and from an Infobar you can access arbitrary external services. This example shows it well:

This one shows an OpenTable reservation link and a Yelp review, but almost anything can go there as long as you can squeeze it into the limited space you have.

The SearchMonkey is still in its infancy. It needs developer support. If you are in Silicon Valley, please come to the Developer Launch Party next week on Thursday May 15. See the link for details. If you aren't in the area, or even if you are, sign up for a developer account at <http://developer.yahoo.com/searchmonkey/preview.html> and help encourage the Web to become more semantic.

Posted by Rasmus in PHP at 14:50

Saturday, January 12, 2008

Looking for a new 1U box

This server is well over 5 years old and really starting to show its age. It's a 2.66GHz P4 with a Gig of ram and mirrored 80G drives.

I am constantly running out of disk space and spamd chews up all available cpu on a regular basis. Moving Bayes token expiration to a cron job helped spamd cpu usage a little bit, but it still isn't a happy server.

I need way more disk and way more ram. I'm looking for a dual-cpu, mirrored (software raid this time) 500G drives, and 4G of ram 1U box to stick in the colo. This <http://www.siliconmechanics.com/i14740/Quad-Core-Xeon.php> fits the bill, although I find the quad-core cpus a bit silly. I'd be fine with quick dual-core chips, but the bang-for-the-buck on the configs there seem to favour the 4-core chips.

This Supermicro <http://www.xpcgear.com/sc813t50c1.html> looks ok as well.

A Sun X2200 is a possibility, but they seem a bit pricy for what you get and I like the new 45nm Intel chips.

Has anybody run across any interesting new 1U servers?

Posted by Rasmus at 17:23

Sunday, December 16, 2007

New toys: iMac and Sonos

A couple of early Christmas presents for the new house. A new iMac for the kitchen nook and a Sonos system. I guess I haven't bought ram in a while because it was under \$100 to upgrade the iMac to 4G with ram from OWC. The iMac is a perfect fit for the kitchen. The black and aluminum matches the counters and appliances in the kitchen nicely.

I had been eyeing a Mac Mini for ages, but Apple doesn't seem very interested in the Mini and for the price the iMac seemed like a much better deal. Bought it from Amazon and it didn't come with Leopard but Apple has an update program so the upgrade is mostly free. They want \$10 for shipping me the CD. I already have another copy, so it would be nice if they would just give me a serial number.

As nice as the iMac is, the Sonos system is more interesting. I have had various mechanisms for playing music from computers to decent speakers over the years, but they were all inconvenient hacks. I don't want to have to use a computer to control the music, and I definitely don't want to do it via a clunky TV-based interface either. The Sonos with its controller does a great job of taking the hack out of the system. It is basically a wireless meshing modular music system. You put either an amplified box with speakers connected, or an unamplified box connected to an existing stereo in each room and the single controller can then control each zone individually or you can link them all up so all the zones play the same music. Each box also has a Line-In that can be used as a source and played in any other zone

I wanted the iMac to be able to go to sleep without killing the music, so I used a 250G Simpleshare drive I had sitting around. One of these days I need to figure out a real NAS system for the house, but for now 250G is plenty as a Sonos media source. I pointed iTunes at the Simpleshare and copied all the music to the drive, then I pointed the Sonos system at the drive as well and it worked nicely.

Internet radio streaming, Rhapsody and Pandora are all working very nicely. I realize this turned out to be a bit of a boring post since there were no technical hurdles and thus no interesting hacks involved in getting any of this working. But that is afterall why people buy things like iMacs and Sonos systems. If they didn't just work without days of fiddling there are plenty of cheaper options that will let you hack and fiddle for months and in the end you get something that almost sorta mostly works.

Posted by Rasmus in Audio/Video at 09:08

Thursday, February 8, 2007

Pipes

<http://pipes.yahoo.com> is a cool toy, and by toy I mean it in the useful and cant-stop-playing-with-it sense. My first impression when I saw an early version a couple of months ago was, "How the heck did they do that?" I was reading the Javascript source code for quite a while. Once you get beyond the fact that this is a browser-based app doing this without Flash, or Java or any similar cheats, you get down to what the app actually does.

Years ago I wrote this silly little Mashup example:

<http://buzz.progphp.com/?q=4>

It grabs an RSS feed, in this case the top daily search term % movers from <http://buzz.yahoo.com/feeds/buzzoverm.xml> which gives you an indication of what is on the minds of web searchers right now. I took these searches and did a Yahoo Image search and a News search and combined them in that oval interface you see. I had to do a bit of RSS and XML parsing to take these different data sources and combine them. This is what Pipes is all about. It provides a visual environment for manipulating data sources and then provides a number of different ways to get the results and integrate them into other things. Directly in your RSS reader is probably the simplest, but you could also feed it to PHP and do further data manipulation.

A simplified Pipes version of the above takes the same Buzz.yahoo.com RSS feed and does a Flickr search on each search term. The result looks like this:

http://pipes.yahoo.com/pipes/DnudMIO32xGDclu7pRr_og

The point here is not the visual output. It is meant to be fed to something else. Hover over the "Subscribe" link on the right there. Then click on the "How this pipe was made" image on the left to see how it works.

This is a particularly lame and simple pipe. Some much cooler ones include:

Blog Buzz for Pipes combines a couple of different blog watching feeds, filters out duplicates and gives you a combined feed in reverse chronological order. When you look at how it was made it becomes immediately obvious what it does. You can save a copy and make your own version that watches for whatever terms you want.

Another interesting one takes the New York Times front page, runs a content analysis on it to get a set of representative keywords and then does a Flickr search on each of those. <http://pipes.yahoo.com/pipes/vvW1cD212xGMiR9aqu5IkA/>.

Here is a much more complex pipe that takes some user input. It finds apartments near things. In this case it looks for apartments within 2 miles of a Park in Palo Alto, California by searching Craigslist, then doing a location extraction and then doing a Yahoo! Local Search for that location.

Even if you have no use for processing data sources this way, open up one of these Pipes and drag the boxes around and watch the pipes react. Web apps don't get any cooler than this right now.

Posted by Rasmus in Software at 02:03

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Tuesday, January 30, 2007

Want a PHP job?

Want to work on some of the busiest and coolest web apps in the world?

Do you like Flickr, and want to work downtown San Francisco?

Or perhaps you are into music, movies or TV and want to work out of Santa Monica? Jumpcut? Or have you seen answers.yahoo.com? Address Book, Personals, Search, Premium Services, Hot Jobs? Want to do interesting things combining PHP and Flash?

Yes, I get a referral bonus, but I need more toys. You get a cool job though, so I think we are even.

Send me your resume and let me know what sort of stuff you are interested in or poke around on <http://careers.yahoo.com/> and let me know which job interests you and I will forward your resume to the appropriate hiring manager.

[edited to remove RSS ad test I had forgotten about]

Posted by Rasmus in PHP at 23:36

Saturday, January 27, 2007

100 Runs

My previous entry way back on Aug.18 talked about the Nike+ipod widget I had picked up along with a first gen Nano. I figure it is time for an update. I still use it 5 or 6 times a week and I just logged my 100th run. There have been a few more actually, but I have had the odd Nano crash/corruption on me before I had a chance to upload a run. Despite those occasional technical glitches, I still really like using it. I have been eyeing the Garmin Forerunner 305, but I actually like listening to music or a podcast as I run and it seems a bit much to strap a GPS to my wrist and also carry an mp3 player. Plus the thing costs \$300+ vs. the \$27 the Nike+ipod gadget costs, assuming you already have a Nano. My 100th run looked like this:

620 miles later and some 80,000 calories burned I have gone from being a 205lb fat slow geek to a 175lb slow geek. Not entirely true, I lost about 10lb before I started running more seriously, so the Aug.18 to Jan.27 timeframe was about a 20lb drop. 80,000 calories and I read that one pound is about 3500 calories which seems to fit my numbers.

I ran a couple of hundred miles in a pair of Nike+ shoes. It is very convenient to have the sensor right in the sole of the shoe, but that is the only thing I liked about those shoes. Now I have a pair of New Balance 767's and have sewn the little sensor seated in a Switcheasy RunAway thing to the front of the shoe with the back looped through the bottom of the laces. That gets it as horizontal as possible and I find it is more accurate that way. Of course the RunAway plastic thing needs a bit of tape to make sure it doesn't swing or pop open. Here is a picture of the hack job:

I have also been playing with the run data a bit. Each run is stored in an XML file which is uploaded to the Nike site by itunes when you synch. You can easily read the files yourself by mounting the ipod and navigating to it. However, the red Nike interface you see above is a flash thing and since flash is obviously client-side the site needs some sort of API. With a little digging with Firebug it was trivial to figure out how it worked and I whipped up a little Serendipity plugin to show my stats. You can see it in the right column of this page. And I wrote a simple PHP 5 class that makes it rather easy to integrate this stuff into other apps. You can see the code here:

<http://lerdorf.com/php/nikeplus.phps>

Each run in the run list has an id and you can call the run() method in that class to get the raw data for that run. Graphing that gives us something like this for that same run I showed at the top:

The grey line is the raw data and the red line is a Lowess curve based on that. I find it interesting how noisy the raw data actually is. Especially for the shorter runs, but by applying a bit of local regression it cleans up nicely and data that seems accurate emerges. I can pick out the uphill stretches of my run from the graph above. If you click on the runs in the right sidebar section there you can see the graphs of my other runs.

If I could somehow find some free time I'd love to build a better tracking site than that Nike thing and then also have it support the Forerunner and whatever other devices are out there. But for now most of my free time is spent running. Probably better for me in the "long run" anyway.

Posted by Rasmus at 17:21

Friday, August 18, 2006

Nike+iPod Goodness

On a whim I picked up the Nike+iPod doodad the other day and this morning did my first decent run with it and I am quite impressed. I'm not much of an Apple nor Nike fan. I do have a Powerbook, but doubt I would buy another one (see previous Toys entry), but this little device is definitely cool. It is going to make me run further and faster and at a more consistent pace. If you already have a Nano you really should pick one up. It's only \$29, \$27 if you work somewhere where you get the Employee discount. Even if you aren't a runner, this works perfectly well for walking as well. Even if you don't have a Nano, you can pick one up for \$108 in the refurbished section of the Apple store right now. Refurbished Apple items seem to be indistinguishable from new items in my experience.

It comes as two tiny devices. Doodad1 and Doodad2. Doodad1 you attach to your Nano as pictured on the right. It only comes in white so it doesn't match my black Nano all that well. Not that I care. Doodad2 you attach securely to your shoe. It should be on tightly such that it moves with your foot and doesn't dangle in a pouch or something. You can of course buy the \$100 Nike+ shoes which I am sure is the goal of this, but I find it works perfectly well stuffed under my sock on top of my foot with the laces tight below and above it. I don't feel it and it seems to be very accurate measuring distances. It uses some sort of piezoelectric accelerometer, like in a speaker, to measure the amount of time your foot spends on the ground combined with the time between foot strikes to figure out how fast and how far you are running. After I calibrated mine over a 1 mile distance it seems to be very accurate. I tried running 400 meters with long strides and then the same 400m with short strides and it didn't get confused. Walking the same 400m it managed to measure accurately as well. EEtimes had a good article on it.

The guts of the Doodads look like the image below.

While you are running the screen shows your progress. Distance, pace and time. If you hit the center button it will use either a male or a female voice to tell you the same information so you don't need to look at your screen. If you choose a set distance for your workout it will tell you that you are "halfway", "400 meters to go", "300 meters to go"... I thought it was interesting that it used meters even though I have mine set to give me everything in miles. I go both ways on the miles/meters thing so I don't care, but I could see someone being confused by that. When you finish your workout and you have beaten your previous best time or distance record you get a little congratulatory voice by Lance Armstrong or a couple of others, I think. I have only gotten Lance so far.

You then plug your iPod back into your computer and it uploads (if you let it) your workout to the Nike+ site. This site has a very nice Flash app (for some definition of nice tempered by the fact that it is Flash) which keeps track of your runs. It's a very shiny app that shows each individual run with time, distance, pace and calories burned (you enter your weight during setup).

The yellow line shows your speed during a run. The dots along the run there appear to be the times I pressed the center button to hear my progress. When you mouse over them in the app you see the distance and pace at that point. The speed seems to match my mental state and the hills along the run pretty well. And I find it really nice to be able to quickly check if I am falling behind the pace I know I can run. It is too easy to trick yourself into believing you are running at your optimal pace.

There is an overview screen where you see your runs. I have only done 2 so far. A short 1 mile run to check the calibration and then this morning's 10km run. When you mouse over the bars it gives you the details on each run.

And there is a summary screen showing your farthest run, best 1mi, 5km and 10km runs.

There are also ways to set goals and to create groups where I assume you can see other peoples' runs and motivate each other to run more that way. I don't know anybody else with one of these yet, so I haven't played with that feature. Perhaps we need to set up a running group for fat grumpy open source developers. It might help turn us into just grumpy open source developers. My wife showed an immediate interest in it as well, but it doesn't support multiple

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

profiles on the same sensor. You can however have multiple sensors and share the same iPod, so I already ordered another \$27 sensor for her.

Posted by Rasmus in Audio/Video at 09:25

Sunday, August 13, 2006

Time to buy a Thinkpad again?

About 18 months ago I had a very nice perfectly configured Debian T42p Thinkpad stolen at a conference. I had less than two weeks before the next conference and I didn't have time to fiddle with my OS to get a new Thinkpad up and running with Debian. It also annoyed me that there was no way to buy a new Thinkpad without paying Microsoft for yet another Windows XP license I would never use. So I swallowed some Open Source pride and bought a Powerbook G4. It took a bit of getting used to, but overall it was a pleasant experience for a while. It was a Unix laptop which just worked and I was up and running quickly.

The Powerbook hasn't done so well over time though. It is not nearly robust enough for my hectic travel schedule which include plenty of mad dashes through airports in strange countries. The clasp to hold it closed broke long ago. The power plug thing is so bent out of shape it won't hold the power connector solidly. These two factors probably contributed to my battery completely dying after about a year and I had to replace it. Also, whenever I put it down the CD Rom tries to eject a non-existent disk, and the screen has a long dent in it which causes a dark area on the display. The cheap silver paint has rubbed off the case where my wrist rests, the key labels are disappearing and my cursor-left key isn't there anymore. Half the time it won't come out of sleep mode requiring a reboot and the other half it simply turns itself off. It also gets amazingly hot. I do have full Appicare on it, but it doesn't appear like they will fix any of this stuff because the 4 or 5 big dents in it is obvious signs of it being dropped and they state they won't fix damage due to a drop.

I have of course been eyeing the new MacBookPro, but I see no signs that it would be any more robust. And even worse, now with their Intel move they are using Intel's Trusted Computing Platform for who knows what. I haven't seen much on what they are doing with the TPM beyond tying the OS to the hardware, but since they are not a member of the Trusted Computing Group there is no oversight. I also find myself using very little Apple software. It is shiny and pretty for the most part, but not really something an old UNIX hack like me would use on a day to day basis. Compiling things on a G4 is painful. gcc and gdb do work, but Valgrind and Callgrind don't which means I can't do any serious development on it.

So, back to the Thinkpad. With Lenovo's announcement that they will sell the T60p with Suse pre-installed I am assuming this means I can finally buy a Thinkpad without paying the Microsoft tax. Combine that with the amazing improvements of Linux on the desktop spearheaded by Ubuntu and it all means that my Apple experiment is coming to an end after about 18 months.

Of course now I need to figure out how to navigate the Lenovo ordering process. It doesn't appear that they have this non-Windows T60p available yet. Probably need to wait until after Linuxworld next week. Will also need to figure out whether EPP (employee) or SPP (stockholder) pricing is cheaper, or if there is some other mechanism to get a bit of a discount. How about an AOSD (Annoying Open Source Developer) discount?

Posted by Rasmus at 11:18

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Sunday, March 5, 2006

Get your geo plugin here

PHP 5, JSON from pecl cvs, Map Tile API and the Yui Connection Manager.

Toss them in a bag and shake and you get something like this. Enter a city name, or even a pseudo-name like Philly or SFO and hit return. Each time you enter a new name you will get a map tile for that city.

It's not all that fancy, but it sure is easy to do:

33 lines total which includes about 6 lines of PHP which consists mostly of building the URL to the map tile service. Then the Javascript which has a callback function to read the form values and make the backend GET request and a simple function to take the JSON response and add the map tile to the document. And finally the actual HTML form.

With a slight tweak you can change it to make a nice geocode lookup field.

So, fire up your text editors and start writing some plugins for blog and forum packages and perhaps image gallery applications as well.

Posted by Rasmus in PHP at 15:50

Monday, February 27, 2006

The no-framework PHP MVC framework

March 1, 2006 - Disclaimer: Since a lot of people seem to me misunderstanding this article. It isn't about OOP vs. Procedural programming styles. I happen to lean more towards procedural, but could easily have gone more OOP. I simplified the code a bit for brevity, but have added a light OO layer back in the model now. Not that it makes a difference. What I was hoping to get across here is a simple example of how you can use PHP as-is, without additional complex external layers, to apply an MVC approach with clean and simple views and still have all the goodness of fancy Web 2.0 features. If you think I am out to personally offend you and your favourite framework, then you have the wrong idea. I just happen find most of them too complex for my needs and this is a proposed alternative. If you have found a framework that works for you, great.

So you want to build the next fancy Web 2.0 site? You'll need some gear. Most likely in the form of a big complex MVC framework with plenty of layers that abstracts away your database, your HTML, your Javascript and in the end your application itself. If it is a really good framework it will provide a dozen things you'll never need.

I am obviously not a fan of such frameworks. I like stuff I can understand in an instant. Both because it lets me be productive right away and because 6 months from now when I come back to fix something, again I will only need an instant to figure out what is going on. So, here is my current approach to building rich web applications. The main pieces are:

PHP 5
Yahoo! User Interface Library
JSON

MVC?

I don't have much of a problem with MVC itself. It's the framework baggage that usually comes along with it that I avoid. Parts of frameworks can be useful as long as you can separate the parts out that you need. As for MVC, if you use it carefully, it can be useful in a web application. Just make sure you avoid the temptation of creating a single monolithic controller. A web application by its very nature is a series of small discrete requests. If you send all of your requests through a single controller on a single machine you have just defeated this very important architecture. Discreteness gives you scalability and modularity. You can break large problems up into a series of very small and modular solutions and you can deploy these across as many servers as you like. You need to tie them together to some extent most likely through some backend datastore, but keep them as separate as possible. This means you want your views and controllers very close to each other and you want to keep your controllers as small as possible.

Goals for this approach

Clean and simple design

HTML should look like HTML

Keep the PHP code in the views extremely simple: function calls, simple loops and variable substitutions should be all you need

Secure

Input validation using `pecl/filter` as a data firewall

When possible, avoid layers and other complexities to make code easier to audit

Fast

Avoid `include_once` and `require_once`

Use APC and `apc_store/apc_fetch` for caching data that rarely changes

Stay with procedural style unless something is truly an object

Avoid locks at all costs

Example Application

Here is the example application I will be describing.

It is a form entry page with a bit of Javascript magic along with an sqlite backend. Click around a bit. Try to add an entry, then modify it. You will see the server->client JSON traffic displayed at the bottom for debug purposes.

The Code

This is the code layout. It uses AJAX (with JSON instead of XML over the wire) for data validation. It also uses a couple of components from the Yahoo! user interface library and PHP's PDO mechanism in the model.

The presentation layer is above the line and the business logic below. In this simple example I have just one view, represented by the add.html file. It is actually called add.php on the live server, but I was too lazy to update the diagram and it really doesn't matter. The controller for that view is called add_c.inc. I tend to name files that the user loads directly as something.html or something.php and included files as something.inc. The rest of the files in the presentation layer are common files that all views in my application would share.

ui.inc has the common user interface components, common.js contains Javascript helper functions that mostly call into the presentation platform libraries, and styles.css provides the stylesheet.

A common db.inc file implements the model. I tend to use separate include files for each table in my database. In this case there is a just single table called "items", so I have a single items.inc file.

Input Filtering

You will notice a distinct lack of input filtering yet if you try to inject any sort of XSS it won't work. This is because I am using the pecl/filter extension to automatically sanitize all user data for me.

View - add.html

Let's start with the View in add.html:

The main thing to note here is that the majority of this file is very basic HTML. No styles, or javascript and no complicated PHP. It contains only simple presentation-level PHP logic. A modulus operation toggles the colours for the rows of items, and a loop around a heredoc (

Posted by Rasmus in PHP at 14:39

Sunday, February 26. 2006

Apache 1.3.34 Debian Package w/ mod_deflate and no pthreads

Debian's Apache1 package doesn't quite do what I need. I have been building my own and overwriting the files from the Debian package, but that can get annoying. So I hacked my changes into the Debian source package and built real .debs. I figure they might be useful to others.

The main changes are to get rid of -lpthread (needed by mod_perl) and -lexpat and to add mod_deflate. To enable mod_deflate make sure your /etc/apache/modules.conf file has:

```
AddModule mod_deflate.c
```

And in your httpd.conf add:

```
DeflateEnable      on
DeflateMinLength   1024
DeflateCompLevel   8
DeflateProxied     on
DeflateDisableRange "MSIE 4."
DeflateVary        on
DeflateTypes       text/css
DeflateTypes       text/plain
DeflateTypes       text/rtf
DeflateTypes       text/xml
DeflateTypes       text/javascript
DeflateTypes       image/vnd.dwg
DeflateTypes       image/vnd.dxf
DeflateTypes       application/msword
DeflateTypes       application/vnd.hp-HPGL
DeflateTypes       application/vnd.ms-access
DeflateTypes       application/vnd.ms-excel
DeflateTypes       application/vnd.ms-powerpoint
DeflateTypes       application/vnd.ms-project
DeflateTypes       application/vnd.visio
DeflateTypes       application/x-javascript
```

Posted by Rasmus in Software at 23:06

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Tuesday, January 3, 2006

Updating my road warrior kit

My wife works for APC, so APC stuff is readily available. Here are a few things that I am looking at for my road warrior kit.

I'd like to get to the point where I just have 1 plug and I am almost there. The TravelPower adapter plugs into the wall/car/plane and it then powers any USB devices like the mobile wifi router and the iShuffle (if I had a cellphone - I hate cell phones, it would power that too), and then the main power feeds into the universal battery which in turn powers the Powerbook. When I unplug from the wall the battery takes over and combined with the internal Powerbook battery I have about 130 Watt-hours, which for me translates into about 7 hours of laptop use. This is what it looks like:

The above description is what I'd like it to do, but unfortunately it doesn't quite work. The stumbling block is the TravelPower adapter thing.

My first big beef with it is that it weighs a ton. My second, and more serious problem, is that it only goes to 20V and not 24V like the universal battery. 24V happens to be what the Powerbook needs. I could perhaps live with that if I could feed 20V to the universal battery and have it output 24V to the Powerbook, but that doesn't work either. I can charge the battery with it, and I can then power the Powerbook with the battery, but I can't do both at the same time which defeats my goal of being able to just plug in my bag in a single plug in the airport. My smaller and much lighter Targus travel adapter is fine for plane/train/car power, but since it doesn't plug into a regular plug as well, I still need to bring something else. For now it will just have to be my regular Powerbook power brick which can charge the universal battery and the Powerbook at the same time and gets me most of the way to the single plug in the airport, it just means I have to switch to the Targus in the plane if I have power there. My product suggestion to APC is a new version of the TravelPower adapter which is half the size and a quarter of the weight and can go to 24V.

Next, the universal battery, or UPB80, is nice. It has been with me around the world a few times now. It is small, light and cheap. You get 80 Watt-hours for about \$150 when a 50 Watt-hour internal Powerbook battery costs about the same.

It can be a little bit flaky. If you don't set the voltage and plug things in in the right order, it doesn't work. To fix it, just unplug the main connector from it and plug it back in and it goes. It supports any laptop I can think of with its 15V, 16V, 18V, 19V, 20V and 24V settings along with a bag of various tips. You plug your existing laptop power into one end of this silver cable that comes with it, plug the middle into the battery and the other end into the laptop to charge the battery and power your laptop at the same time. Then when you disconnect from the wall the battery kicks in. Don't forget that as far as your laptop is concerned it is still running in powered mode. It has no idea you are using an external battery, so when running from this thing you have to manually switch your laptop into low power consumption mode for maximum battery life. Not much else to say about this one. It works. I don't leave home without it.

I also recently picked up this very cool universal plug thing that supports all the plug types in the world in a very small form factor. Hard to describe. It looks like this (unfolded):

Try clicking on the above picture. It shows the various ways it can unfold itself Optimus Prime-style.

The latest addition is the WMR1000G Wireless Mobile Router. This thing can do AP, Router and Client mode just like my hacked WRT54G, but it is tiny. I probably still won't give up the WRT since I use it to hook up my huge antenna for the

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

extreme situations where I am far from an access point, but even then, I can put it back to back with the WRT and rebroadcast the signal locally. I was doing that through the Powerbook before, but it meant I had to be tethered to the WRT. Now I can be wireless too and still use the big antenna to pick up remote networks.

It is interesting that it comes with a USB power cable along with a regular small power block. I tried powering it from my Powerbook USB port, but it didn't work, and there is also a big warning in the instructions that it is only designed to be plugged into the TravelPower Adapter thing (which doesn't help me much since that silly thing doesn't support my Powerbook).

There is a little switch on the side that has 4 settings. AP, AP/Router, Config and Client. To configure it, stick it in Config mode and connect either via the wire or the wireless. ESSID is "default" in Config mode, "APC_Router" in AP/Router mode, and "APC_AP" in AP mode. The difference between AP and AP/Router mode seems to just be that in Router mode you get DHCP/NAT, whereas in AP mode it is just passthrough. In Config mode (login using Admin/APC) you have 4 main configuration screens named, System Setup, AP Mode, AP/Router Mode, and Client Mode. The config is pretty simple. System Setup has:

So you can set wired and wireless MAC addrs that can configure the thing without needing a passport. Useful for the folks smart enough to figure that out but dumb enough to forget their password?

On the AP Config screen you have:

No surprises here either. "Trusted Stations" means MAC filtering, and it supports WEP and WPA-PSK.

On the AP/Router config screen, things get a bit more interesting.

I am not sure what "Travel Mode (Hotel)" means under Connection type there, but the other options are PPPoE, PPTP, L2TP and Static IP so I guess Travel Mode means DHCP here. Weird. And under Advanced we see this odd-looking menu:

I find it odd to see "Age of Empires" there. The last and only option missing from the screenshot view is "Yahoo! Messenger", so no other games on the list. The next menu over, "Port Forwarding" lets you fully configure it, thankfully.

The DDNS screen has support for DynDNS, DtDNS and Cn99. The Network Diag screen lets you do pings and dns lookups. Under Options you can configure backup DNS servers and enabled UPnP (enabled by default). The PC Database screen lets you hardcode static ips for the dhcp server - Nice! And the security screen looks like this:

Not entirely sure what they mean by a DoS firewall. What sort of DoS attacks is it firewalling?

And finally the Client mode config screen.

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

This is the first place I think this router come up a bit short. There doesn't appear to be any way to have it discover available access points. The one you see on the screen there is one I configured manually for my AP at home. Once you configure it to connect to a specific wireless network, it works fine and basically just turns a wired connection into a wireless one. I use this all the time with my WRT on the road because the radio is stronger and more sensitive than my Powerbook's wifi and with the big antenna connected there is no comparison.

Without the ability to connect an external antenna and the lack of an AP Scan feature in Client mode, I don't see it replacing my WRT, but I still really like the features this tiny little thing has. And because of its size and weight it has definitely earned itself a spot in my travel bag.

Posted by Rasmus in WIFI Toys at 23:17

Sunday, November 6, 2005

More fun with the Yahoo! Maps API

I spent a few hours with the Javascript-Flash Yahoo! Maps API today. The result is here:

<http://lerdorf.com/map>

Have a look at the source. There is a source link at the top-right on the page. It is about 90 lines of HTML and Javascript and virtually no server-side scripting.

I usually throw PHP code at all sorts of problems, but in this case I could do pretty much everything I wanted to directly from Javascript via the excellent API. The initial zooming from way out is a bit distracting and doesn't work too well, but I wanted to play with that a bit. Once you are zoomed in and searching for things, the search is updated as you move around the map and everything is event-based so the page doesn't need to be redrawn from scratch.

Both input fields are free-form. You can put a city name, a zip code or a full address in the Location field and in the What field you put what you are looking for. There is a special hack that checks for something like 4* and the filters the results to only show you those entries that were rated 4* or higher on Yahoo! Local. You can of course also just put something like "pizza" or "mexican" in that field.

Most of the magic here is done by 2 things. The LocalSearchOverlay and the event handling. Note how Map.EVENT_MOVE and Map.EVENT_ZOOM_END are registered events in the onInitialize() function. When you scroll the map or zoom it the onOverlayInit function will get called and the LocalSearchOverlay will be recalculated for the new map coordinates. Same thing happens when you change something in the input fields. The updateMap() function is called which will center the map at the new location and update the LocalSearchOverlay appropriately. There are a few more Javascript tricks here and there in it, like updating the link at the top so you always have a way to grab a link to your current search and send it to someone, but other than that there really isn't all that much to figure out here. Once you understand which events happen when and which methods are available where, you can do some really powerful things with this. It is all documented here:

<http://developer.yahoo.net/maps/flash/V2/flashReference.html>

and people are discussing it here:

<http://groups.yahoo.com/group/yws-maps/messages>

In my last entry I showed how to parse a geocoded XML file and put markers for each entry on the map. Someone asked me how I would do the using the JS-DHTML API instead of the JS-Flash API. It's a whole lot harder in DHTML, but it works. You can see it here:

<http://lerdorf.com/php/ymap/dquakes.php>

And I talk a bit about it here:

<http://groups.yahoo.com/group/yws-maps/message/612>

Posted by Rasmus in Software at 17:09

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Thursday, November 3, 2005

GeoCool!

Web 2.0 and the programmable web that I and others have been talking about for a while has mostly been vapourware so far. There are a few generic components that are useful, but it is somewhat limited what you can do with them. And yes, you may consider this a somewhat biased view, but I think Yahoo!'s new geocoding platform is a huge step in the right direction.

There is of course the fancy new maps.yahoo.com/beta site which is fun, but as far as I am concerned the killer app here is the geocoding platform that drives this. And it is completely accessible for anyone to use. It's also a sane API that anybody can figure out in minutes. Here are a few tips for using this API from PHP 5.

Step 0 - The raw geocoding API

Whenever I do anything with web services, I always add a request caching layer. So here are the base building blocks implemented in 2 functions. One for doing request caching and the second to do the actual REST query to the geocoding service.

```
Result['precision']; foreach($xml->Result->children() as $key=>$val) { if(strlen($val)) $ret[(string)$key] = (string)$val; } return $ret;?>
```

The above code is the contents of `geo.inc` which you will see included in the following examples.

Easy enough? No real tricks here. We simply send a regular GET request to <http://api.local.yahoo.com/MapsService/V1/geocode> with the location parameter set to an address. You can try it yourself directly from your browser by clicking here:

<http://api.local.yahoo.com/MapsService/V1/geocode?appid=rlerdorf&location=701%20First%20Ave,%2094089>

You can read more about the geocoding service here:
<http://developer.yahoo.net/maps/rest/V1/geocode.html>

Step 1 - Writing your first application

We can just toss a form around this and dump the results to make sure things are working.

GeoCoding API Example

You can see this one in action here:

<http://lerdorf.com/php/ymap/geo1.php>

Note how it is able to fill in missing details for a partial address. eg.

<http://lerdorf.com/php/ymap/geo1.php?location=701+First+Avenue+94089>
results in:

```
[precision] => address  
[Latitude] => 37.416384  
[Longitude] => -122.024853  
[Address] => 701 FIRST AVE  
[City] => SUNNYVALE  
[State] => CA  
[Zip] => 94089-1019
```

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

[Country] => US

This means that you can use it for a bunch of different things. Address to lat/long, of course, but also address to city, or city to zip code conversions. Or 5-digit zip to 5+4. This is of course rather US-centric right now, but that will improve over time.

Step 2 - Adding a map

The geocoding is cool, but an actual map is cooler. Easy enough:

```
#mapContainer { height: 600px; width: 800px; } GeoCoding API Example mymap.addMarkerByLatLon( new CustomP  
OIMarker("  
", "", '0x0012f0', '0xFFFFF'), new LatLon());
```

You can also let the API figure out your markers for you which makes this even simpler. If the RSS feed is using georss correctly you can use the GeoRSSOverlay mechanism. Here it is using the earthquake RSS feed directly:

<http://lerdorf.com/php/ymap/rssquakes.php>

And here is the code. I am still loading the RSS feed myself from PHP because I want to get the pubDate and title from it, but everything else is handled automatically.

```
#mapContainer { height: 600px; width: 800px; }
```

Posted by Rasmus in Software at 15:32

Thursday, September 1, 2005

Flickr API Fun

I like stuff I can pick up and do something useful with in an hour or two. Perhaps my attention span is too short, but if I have to read a 300 page spec before I get to Hello World, then it's not for me. Or you would at least have to pay me a lot of money to suffer through it. I think people refer to this as "immediacy". For me I think it is mostly lazyness. If I can't figure it out in an hour, it's broken as far as I am concerned.

Flickr's REST API is not broken. You can read all about it at <http://flickr.com/services/api>.

There are links there to various wrappers for the API, but I ended up writing my own. I have a bad habit of doing that. This entry will focus on my PHP wrapper for the Flickr API. It is based on Cal's version and is compatible with it, but it expands on it and puts some PHP 5.1 features to good use. You can see it here:

http://lerdorf.com/php/flickr_api.phps

Before you get started, in case you want to follow along, go get yourself an API key at

<http://flickr.com/services/api/key.gne>

You will need two pieces of information to fully use the API. An API key and an API secret. And if you are going to do anything that requires authentication, you need to set a callback url as well. More on that later. To get your secret after applying for and getting your API key, go to

http://www.flickr.com/services/api/registered_keys.gne and click on "Edit Configuration".

Many functions in the API do not require authentication. Getting a list of someone's public photos, for example, is something anybody can do by just browsing Flickr, or by just going to this URL:

http://flickr.com/services/rest/?method=flickr.people.getPublicPhotos&user_id=56053642@N00&api_key=3aba8184848f9263b80795c95529bcd1

Guess what, you just sent a REST Web Services query.

Or, slightly cooler. A list of tags related to the tag you provide based on Flickr's clustering code.

http://flickr.com/services/rest/?method=flickr.tags.getRelated&tag=monkey&api_key=3aba8184848f9263b80795c95529bcd1

The whole point of web services is to provide data in a machine-readable way so you can do something more interesting with it. That's where the API wrapper comes in. You can of course also use Flickr's feed mechanism to do this.

But back to the API and the PHP wrapper. Getting a list of someone's public photos is done like this:

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

```
$secrets = array('api_key'=>'your_key_here','api_secret'=>'your_secret'); $flickr = new Flickr($secrets); $photos = $flickr->peopleGetPublicPhotos('56053642@N00');
```

This will give you an array of photos. Or to be precise, an array of information about the photos. Note the mapping of the method name. `$flickr->peopleGetPublicPhotos` maps to `flickr.people.getPublicPhotos` in the documentation. And the returned XML is converted to a more useful (and more memory-cacheable - I'll write something up soon on that) PHP array. The example result XML for 2 photos looks like this:

Which gets mapped to this PHP array (in `print_r` format):

```
Array (
  [page] => 1
  [pages] => 1
  [perpage] => 100
  [total] => 2
  [photos] => Array (
    [39006009] => Array (
      [id] => 39006009
      [owner] => 56053642@N00
      [secret] => f2086066d5
      [server] => 33
      [title] => IMG_7564.JPG
      [ispublic] => 1
      [isfriend] => 0
      [isfamily] => 0
    )

    [39006000] => Array (
      [id] => 39006000
      [owner] => 56053642@N00
      [secret] => 4ec57bd51f
      [server] => 28
      [title] => IMG_7551.JPG
      [ispublic] => 1
      [isfriend] => 0
      [isfamily] => 0
    )
  )
)
```

To turn a photo into a URL you can use in an IMG tag you would call the `$flickr->getPhotoURL()` method. It isn't very complex. Here is what it does:

```
function getPhotoURL($p, $size='s', $ext='jpg') { return "http://photos{$p['server']}.flickr.com/{$p['id']}{$p['secret']}{$size}.$ext"; }
```

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

The default size is the small 75x75 square thumbnail. See the URL Documentation for further info. So here is the full code to put up the first 50 thumbnails from someone's photostream:

The contents of secrets.inc is just that \$secrets array I referred to above. You can see the output of this script at flickr_demo1.php.

Flickr users are uniquely identified by a very cryptic-looking nsid. You don't see this id anywhere when you are clicking around on Flickr. But you can look up a user's nsid if you know their photo url, their user name or their email address. flickr_demo2.php shows you how to do that. Change the u= parameter in the URL to look up other users.

Playing with the non-authenticated functions of the API can get you far, but Flickr also lets you authenticate, and it will let the users using your application authenticate themselves. That lets you do a whole class of cool things that something like the RSS feed mechanism doesn't provide. For example, I wrote a Gallery to Flickr migration tool that can take my Gallery photo albums and copy the pictures to Flickr and put them in a Flickr set with the same name as the Gallery album they came from. You could also write alternative frontends for it and integrate your Flickr photos with your own web site. Or perhaps write a Gallery plugin that uses Flickr as the backend. All sorts of possibilities here.

But in order to do any sort of reading of non-public information or writing to your Flickr account via the API, you have to authenticate. Flickr uses a token-based authentication system where you make a roundtrip to flickr.com for the user to log into his flickr account and choose whether or not to grant your application the requested level of access. That means that your application never sees the user's credentials, but instead gets a token with the appropriate rights associated with it that it can then use. Each API call then includes this token, the application's key and all the arguments for whatever method you are calling and a signature using your application-specific secret across all the arguments. That means that even if someone sniffs your traffic, all they can do is replay the exact API call. They can't use it to execute arbitrary things against your account. Users can also remove an application's access later by going to <http://www.flickr.com/services/auth/list.gne>. The system is described at <http://www.flickr.com/services/api/auth.spec.html> and is worth a read if you are interested, but you don't really need to understand it. Just use the wrapper. Here is how:

This probably looks a bit cryptic. This says that if you already have a token, we simply create a \$flickr object and we are ready to go. If there is no token on the request and there is no 'frob', then redirect the user to the authentication URL which is generated by the call to \$flickr->getAuthUrl with the desired permission level as an argument. The user will then get sent back to your callback url, which you would set to this same script most likely, and on that callback we still don't have a token, but you will be called with a frob parameter. A call to \$flickr->getFrobToken turns the frob into a token. You actually get back an auth array containing not just the token but also the user's nsid, permission level for the token, username and fullname. The idea is then that you include the above

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

blurb on your pages, as flickr_auth.inc, for example and pass the token along from page to page in your web application.

Now we can write our first full little authenticated example. Not much to it. We just call `$flickr->authCheckToken` on our token to see what Flickr thinks of the token we are using.

You can see the output by clicking on flickr_demos.php and selecting flickr_demo3.

So, by having those 3 includes at the top, by the time we get control we have a fully authenticated `$flickr` object that we can start using.

So, an all-out demo. In flickr_demo4 we upload a photo:

```
$photo_id = $flickr->upload($fname,$title,$desc,$tags,$perms,0);
```

Check to see if you already have a set named "Sample Set". If you don't, create it (adding the uploaded photo at the same time):

```
$set = $flickr->photosetsCreate("Sample Set", $photo_id);
```

If you do already have that set, add the uploaded photo to it:

```
$flickr->photosetsAddPhoto($set_id, $photo_id);
```

Then we can add a note:

```
$note_id = $flickr->photosNotesAdd($photo_id,342,70,50,50,"This is Carl");
```

Get info on the photo:

```
$photo = $flickr->photosGetInfo($photo_id);
```

And get a direct URL to it:

```
$url = $flickr->getPhotoURL($photo,'m');
```

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Have a look at the full source code and try it by going to http://lerdorf.com/php/flickr_demos.php and clicking on demo4.

Even if you don't use Flickr, I think there are a lot of interesting things here. An interesting web service authentication mechanism, a nice and clean REST API that allows for complex operations, and some PHP 5.1 XML and stream handling if you look closely at the flickr_api code.

Posted by Rasmus in Software at 09:00

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Wednesday, March 2, 2005

Buzzing the Yahoo! Search Web Services

March 22. Update: And here is the Flickr version: flickr.progphp.com

PHP5 has a revamped XML architecture that makes dealing with SOAP and REST Web Services extremely simple. I wrote a little demo application against Yahoo!'s new search web services. It uses the various search buzz RSS feeds to seed it or you can provide your own search terms. It then uses those terms to pull image, web and news search results which it arranges somewhat haphazardly. You can play with it at <http://buzz.progphp.com>. The orange box shows results from a news search, and when a term doesn't have enough news hits I supplement them with web search results which is shown in green to distinguish them.

Apart from a bunch of messy CSS, the application is actually quite simple. Pulling from the RSS and REST servers is trivial. Here is a one-liner to pull an RSS feed from a url:

```
$url = 'http://buzz.yahoo.com/feeds/buzzoverl.xml';  
$xml = simplexml_load_file($url);
```

The actual implementation wraps this and returns an associative array with just the title and the link, like this:

```
foreach($xml->channel->item as $item) {  
    $ret[(string)$item->title] = (string)$item->link;  
}  
return $ret;
```

For the search REST queries it isn't much harder. You build your query string:

```
$url = 'http://api.search.yahoo.com/';  
$url .= 'ImageSearchService/V1/imageSearch';  
$url .= '?query='.rawurlencode($q);  
$url .= "&appid=$appid";  
$url .= "&results=$results";  
$url .= "&type=$type";
```

I then throw a cacheing layer in front of all these so I don't hit the feeds on every request. The core of the cache layer looks like this:

```
$stream = fopen($url,'r');  
$tmpf = tempnam('/tmp','YWS');  
file_put_contents($tmpf, $stream);  
fclose($stream);  
rename($tmpf, $dest_file);
```

A straight fopen() can be used since this is a simple REST query and the result is streamed directly to a temp file which is then renamed when complete to make sure other processes never see a half-written file. Check the mtime on \$dest_file and read it until it gets too old, then refresh it.

Although I am not using any SOAP in this particular example, it isn't much harder to pull from a SOAP service. Here is a simple example that pulls from Amazon's SOAP service (they have a REST interface as well). It caches a serialized version of the generated object based on the service index and keywords requested.

```
$amazon_index = array(
```

```
'DVD', 'Photo', 'Electronics', 'OfficeProducts', 'HealthPersonalCare',
'Toys', 'Baby', 'VideoGames', 'MusicTracks', 'OutdoorLiving',
'Blended', 'MusicalInstruments', 'Magazines', 'DigitalMusic',
'Jewelry', 'Video', 'Tools', 'PCHardware', 'SportingGoods',
'Classical', 'Software', 'Books', 'VHS', 'Wireless', 'Restaurants',
'Music', 'GourmetFood', 'Miscellaneous', 'Kitchen', 'WirelessAccessories',
'Merchants', 'Beauty', 'Apparel'
);

function amazon($index, $keywords, $timeout=7200) {
    $dest_file = "/tmp/aws_{$_index}_".md5($keywords);
    if(file_exists($dest_file) && filemtime($dest_file) > (time()-$timeout)) {
        $result = unserialize(file_get_contents($dest_file));
    } else {
        $aws = new SoapClient('http://webservices.amazon.com/'.
            'AWSECommerceService/US/AWSECommerceService.wsdl',
            array("trace" => 1));
        $result = $aws->ItemSearch(array(
            'SubscriptionId'=>'XXXXXXXXXXXXXXXXX',
            'AssociateTag'=>'lerdorf-20',
            'Request'=>array(array('SearchIndex'=>$index,
                'Keywords'=>$keywords))
        )
        );
        $tmpf = tempnam('/tmp','YWS');
        file_put_contents($tmpf, serialize($result));
        rename($tmpf, $dest_file);
    }
    return $result;
}
```

I still much prefer the REST services out there. SOAP always reminds me of being stuck behind the guy in a hat driving a Lincoln Towncar. You eventually get to where you want to go, but the journey is painful. With REST you can just toss your query into your browser and have a look at the returned XML. SOAP starts to make more sense when the queries you are sending get more complex than just tossing a couple of keywords to a search service and setting a couple of flags. But don't even try to read the SOAP spec. If you managed to fight your way through that spec already, try the new WSDL 2.0 Draft Spec. This is the sort of stuff that makes my brain hurt.

And yes, I know the thumbnails don't jump to the front in IE. IE's z-index handling on position: absolute elements is braindead. So use Firefox or Safari or some other browser with decent CSS support. Also, you'll need to let the cookie through. It's just a javascript cookie with your window dimensions so I'll know how big to make the oval. And no, it isn't really meant to be useful. Just a bit of fun visual candy.

Posted by Rasmus in Software at 22:38

Monday, February 14, 2005

Best Chocolate Ice Cream Ever!

This isn't a very high-tech toy, but it makes the best chocolate ice cream in the world. It's a relatively inexpensive Cuisinart ICE device which has a fancy bowl with some sort of liquid jelly inside that freezes and holds its cool temperature for a while. And there isn't much to making ice cream. I have made all kinds, but my current favourite is this recipe I put together:

2 whole eggs
1.5 cups of heavy cream
1 cup of butter milk
1 pinch of salt
4 tablespoons of sugar
1 block (8-10 ounces?) of Trader Joe's broken Ghiradelli dark chocolate
1 block (8-10 ounces?) of Trader Joe's broken Ghiradelli milk chocolate
2 drops of almond extract

Heat up the butter milk in a pot and add the chocolate. Keep stirring and make sure it doesn't burn, although the slightly burnt flavour, if you do happen to burn it, adds character to the ice cream. Once melted into a smooth ganosh-like consistency, stir in the sugar, salt and the almond extract and remove from the heat and let it sit for a bit.

In a clean bowl, whip the eggs for about 3 minutes, add in the cream and whip for another minute or 2. Finally pour the warm chocolate mixture into your whipped eggs and cream mixture and put the bowl in your fridge for an hour or two to cool it down. Trying to make ice cream right away with a warm starting point doesn't work very well. You can stick it in the freezer to speed it up a bit too. When it is cool, take your freezer bowl out and put it on the motor thing and pour your mixture into the bowl. Wait 20-25 minutes and you will have something resembling soft-ice. It's very good at this point, but another 3-4 hours in the freezer and it is perfect.

The secret is of course the pound+ of good chocolate or so that went in which is about double the chocolate of any of the suggested recipes I have seen. Not much sugar is needed because I use half milk chocolate. If you are using all dark or bittersweet, you'll need more sugar. I like the combination of the slightly salty butter milk and chocolate. You can use regular milk and/or lighter cream. However, the fat in the cream is what gives you that rich smooth ice cream taste and texture. If you want to add nuts or fruit, do it right at the end of the 25 minutes and just let them stir in for 30 seconds. If you put them in at the start they tend to clog things up. This toy makes excellent frozen drinks and soft sherberts as well.

Posted by Rasmus in Food at 15:03

Friday, February 11. 2005

TiVo ToGo Annoyances

I put in my request weeks ago to get onto the priority list to get the new TivoToGo-enabled software on my Tivos. It finally arrived a couple of days ago. I wanted to write something up on it as soon as possible, but unfortunately it has taken me 3 days to get anywhere with it, and I still don't have a DVD with my shows on it. There were a number of problems.

The Tivo needed a reboot before it would work and since it was in the middle of recording something I didn't want to reboot it. That meant I couldn't do it the first night.

After a reboot the Tivo Desktop software was able to connect to my Tivo. Unfortunately it takes close to an infinite amount of time to transfer stuff. I have 7 episodes of "24" recorded at best quality I want to put on a DVD. That translates to about 15G of data I needed to transfer. I have a USB wired adapter plugged into a WRT54G and then an 802.11g connection to another WRT54G sitting next to a Windows box. This combination got me a transfer rate of about 1G/hour which means transferring my 7 episodes took over 15 hours.

Ok, a day later all the episodes came across. I grabbed the trial copy of Sonic MyDVD Studio 6.1 which is what Tivo suggests we use to create our dvds. I had configured my TiVo Desktop to stick my shows on E:\tivo since I didn't have room on my C drive. Unfortunately there is no provision in MyDVD to change the directory it looks for tv shows in. You can browse to the files and add them manually, but then you don't get the same ui screen that shows which shows are available. A quick register hack fixed that, but you can't expect Joe User to hack his registry for something this simple.

The MyDVD UI for adding shows is horrible. It doesn't show the episode information at all, so I had my list of 7 episodes all titled "24" with no clue as to which was which. So you are probably better off browsing to them and adding them manually than using the button designed for this.

MyDVD has absolutely no TiVo integration. I am sure that I read at some point that the dvds you create would have a Tivo-like navigation menu. No such thing exists in MyDVD. There isn't even a TiVo theme for the top DVD menu. They could at least have added a JPG with a TiVo logo or something. Nothing.

When I finally did start the burn process MyDVD told me it was building the menus. I had read it was slow, so I went to bed. 7 hours later it was still building the menus. So I left it another 12 hours. No progress. The thing isn't hung. There is a cancel button that works fine and asks me if I am sure I want to cancel since cancelling could take a "long" time!

So as far as I can tell burning a DVD simply isn't working for me. Who knows what the problem may be. Conflict with something else that is installed perhaps? This is Windows, you can never really be sure why something doesn't work. Unfortunately the OSX TiVo Desktop 2.0 is not available yet so I can't do this on my Powerbook yet. At this point it looks like my only real option is to work around MyDVD and use something like this hack to convert the .tivo files to mpeg2. In the long run that's probably much more useful anyway. The .tivo files have a playback password on them you have to remember which seems to me to be DRM run amok. Not that I have a VCR, but if I did, it wouldn't have a password option. This is no different.

Another annoying thing is that the recorded shows have a thin strip of white on black noise at the top. Looks to be too organized to be noise actually. Closed captioning data or something? Guess I may need to edit the stream to cut that out and also to snip out the commercials so I can fit more on a dvd.

I'll update this if I ever get my shows onto a dvd that works. At this point I am having a hard time disagreeing with the TiVo Deathwatch. I'll probably be looking for an alternative soon.

Update: Ok, I finally managed to burn a working DVD from the MyDVD software. After a ctrl-alt-del to kill off the running process that had been going for almost 24 hours, I restarted it, reloaded the project and it actually worked. Still took a couple of hours to sit and chew on the 2 .tivo files I had fed it, but the final DVD works. It still has the noise at the top which is distracting, but the quality is ok. So all in all, out of the 8 hours of "24" which was supposed to be my test case, eliminating the hung MyDVD and other delays and extrapolating from the 3 hours I have managed to burn so far, it will

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

end up taking somewhere between 20 and 24 hours to create the 3 DVDs assuming I am right there at each user-interaction point ready to click or switch dvds and assuming no hung processes or crashes. I think the upper range my patience will tolerate is 8 hours to burn 8 hours of content. For it to be something I would use often it would need to be at least 3x or 4x, as in 15-20 minutes for every hour of content.

By the way, don't get me wrong, I love my TiVo. I have played a little bit with alternatives, and for the core functionality of scheduling and recording shows nothing I have seen beats it. But the alternatives are catching up and my expectations are high.

Posted by Rasmus in Audio/Video at 14:32

Saturday, January 29, 2005

Synching iTunes with rsync

I have an XP box that acts mostly as a uPNP media server for the Philips Streamium and also serves as the server for digital cameras and the other mp3 players. However, since getting the Powerbook and now the iShuffle I have been using iTunes on my Powerbook to manage my mp3's. So the problem was that I wanted to be able to easily update my XP box with the iTunes music library from my Powerbook and let Christine update her iRiver player from it. With the Windows version of iTunes installed and the UMS firmware for Christine's iRiver 390T it is easy to use iTunes to populate the iRiver. The tricky part was synchronizing the libraries including the ratings and play counts. I marked my Windows iTunes folder as shared along with my music folder. I have these on separate drives. If your music folder is in your iTunes directory you could simplify this script a bit:

No iframe support? You can see the code at <http://lerdorf.com/synci.phps>

When you run this it makes your Windows iTunes look exactly like your OSX iTunes by mounting the Windows shares, copying over the hdfm file, munging the paths in the XML file and rsyncing the music folder which means you will lose any library properties local to your Windows box. I would also suggest exiting iTunes on both machines before running this script and synching your clocks from an ntp server. The clock on my XP box drifts badly so I had to add --modify-window to the rsync command.

It should be simple to reverse the script if you want your Windows iTunes to be the master, and with a bit of work you could probably make it go in both directions. But this serves my needs nicely.

Posted by Rasmus in Audio/Video at 10:50

Wednesday, January 26, 2005

iShuffle

A 1G iShuffle showed up today. Carl instantly took to it and at one point declared, "I like it." So if nothing else, Apple has managed to design something that appeals to a 2-year old.

Much has been written about the lack of a screen on this thing and how it is just a USB memory stick with an earphone plug. While that is pretty much true, it is also priced in the ballpark of dumb memory sticks so I am not sure why people complain so much about that. My main use for an mp3 player is during my runs. I tend to run without my glasses so looking at any sort of display while running is out of the question. I have enough trouble picking out cars without my glasses, never mind a tiny lcd display. When I am stationary I turn to my surgically attached laptop which is a perfectly good mp3 player by itself, but running with a laptop with its hard drive spinning would be insane and as far as I am concerned the same applies to running with a hard drive based player. I have already destroyed one HD player by doing that.

We have two other flash players in the household, both of them with screens, and I never look at them. Flash players are not designed to hold your entire music collection. You have that somewhere else and you just load up stuff you want to listen to for a certain activity. Our other flash players have been stuck in random-play mode for years so it was as if Apple designed a player that catered specifically to the way I use them.

Christine has an iRiver which is a great flash player. But they are a tad pricy, and we already have one of those. No fun buying 2 of the same toys. I also have a rather crappy RAVE-MP AMP-256 flash player I picked up from Walmart for under \$50 a while back. It does its job, sort of, but the single AAA battery vs. the AA in the iRiver is annoying, and often it won't start unless I pull the battery and re-insert it. Some sort of firmware bug that they have been promising a fix for. The volume is too low on it and the whole thing feels flimsy, so I am not a huge fan. Then again, for under \$50 it's not a bad deal and if they ever come out with that firmware update that fixes the boot and volume problems, and also allowing you to put more than a 512M add-on SD card in it, then it would be a steal at that price. For now I only see myself using it for its radio and voice recording if I can't steal Christine's iRiver which is better at both.

Anyway, back to the iShuffle. There really isn't that much to say about it. You plug it in, iTunes launches automatically (on both OSX and XP) and you either randomly autofill it or you drag songs onto it. Or you do a combination. The volume is good. It weighs nothing. The controls rival the iRiver stick and are easy to operate with one hand unlike the AMP256 which is impossible to control both tracks and volume with one hand unless you have a triple-jointed thumb. One thing that did annoy me was the fact that it doesn't want to connect to 2 different iTunes. I loaded it up on my Powerbook and then plugged it into an XP box to put a few songs on it that was only on that machine and it asked me if I wanted to tie it to this iTunes instead. I understand why this is so from a legal perspective, but given how trivial this restriction is to work around it is one of those things that is just irritating. Also, when I said no it simply wasn't visible, and it isn't showing up as a regular USB mass storage device after that either. If you don't launch iTunes it appears nicely as a usb mass storage device and you can use it like a regular memory stick. I also really like the built-in USB plug and the fact that it recharges through it. The other two flash players in the household can easily go 6 months between music updates because it is a hassle to find the usb cable and hook it up. With the iShuffle recharging and updating the music on it is a single seamless operation. And with Autofill it is just a single click. The other unique feature I haven't seen mentioned much is the option in iTunes to automatically convert to 128kbps AAC for anything you transfer to it to squeeze more tracks onto it.

Perhaps related to that automatic down-conversion feature, it seems like there may be a bug in iTunes. For example right now it is reporting this:

So which is it? Have I used 984M or 782M? I assume there isn't 200M of overhead hiding in there. I get an out-of-space error when trying to add more tracks at this point, yet they still seem to show up and the 782 number goes up while the 984 stays the same. Unplugging it and restarting iTunes doesn't change anything. I did have a number of large .WAV files that I put on there and then converted them to 128kbps AAC once they were already on there. iTunes is telling me they were converted, but it seems like the fact that the file sizes changed on files already on the player isn't being taken into account somewhere.

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

My final verdict is that this is a cool cheap little player for people like me who can't justify paying \$400 for an mp3 player. The iPod mini is definitely an option here for more features and storage, but there is still a hard drive which won't be happy with my vigorous stomping when I run and up in that pricerange there are other options with built-in radios that seem just as appealing. Since I already have two flash players with radios and the ability to record, I didn't really care about these features, but if you aren't a gadget-collecting geek and just want one and don't mind spending a little bit more I would suggest one of the iRiver flash players. Of course you don't get native iTunes integration with a non-iPod which might be a consideration. All in all I am using the shuffle for my runs now in place of the AMP256 because of its higher capacity and seamless charging/updating and because Christine has claimed the iRiver and won't let me steal it.

Posted by Rasmus in Audio/Video at 23:05

Sunday, January 16. 2005

What's in your travel bag?

I travel quite a bit and a couple of people have asked me what I bring along. So here it is. I emptied my computer knapsack onto the table at the Vancouver hotel I am in after yet another conference. Roughly from left to right:

Spare

wireless cards - Rarely used, but sometimes it is handy to have two

wireless nics in the laptop to bounce a signal along or to lend cards

to people.Charger for the Canon S400 travel camera (camera is taking the picture, of course)

S-Video cable and 8mm to audio left/right cable used for connecting the laptop

to the hotel TV so I can watch movies on a screen bigger than my laptop

screen.USB cable for offloading the camera and also for connecting the flash mp3 player

Case for the camera

ToiletriesTargus Air travel power supplyLightweight extension cordUSB memory stick which doubles as a

penRSA Key to connect to the work network2-foot omni-directional wifi antennaWRT54G + power supply

Spare AAA battery needed for both mp3 player and noise-cancelling headphones

WatchDVI-VGA Adapter cable (critical for presentations)

Credit cards and ATM cardPassport (preferably from a non-war mongering nation)

Laptop + power supplyNetwork cableUSB Mouse (because trackpads suck)mp3 player

Noise cancelling headphones

Nothing really out of the ordinary here except perhaps the rather large antenna. Also, make sure all your power supplies

are 110-240 safe and avoid any surge-protection powerbars (they are too big anyway) since that stuff tends to blow up

when you plug them into a 220V outlet.

Posted by Rasmus in WIFI Toys at 01:10

Monday, December 20, 2004

The Good, The Bad and The Ugly Powerbook

After the rather abrupt loss of my T42p I needed a new laptop quickly. The delay in getting the Thinkpad originally combined with the headache of getting Linux working on it nicely and also to some extent IBM dumping their PC division all contributed to the decision to go with a 15" Powerbook. If IBM would have supported a real OS on their hardware I would have bought another one in a second. I have never been a Mac fan, but I needed a working Unix laptop quickly. After years of having an absolute crap operating system, they finally have a good one. This thing is a 1.5GHz G4 with a half a Gig of ram and the 128M Radeon and then another 512M from Crucial (<http://www.crucial.com>) added on. Would have been nice to put 2G in it, but those 1G sticks are really expensive. After using it for a week and a half I am quite pleased with it overall. Plenty of little gripes, but overall it is a useful little machine.

The Good (with a few reservations)

Stuff works. I didn't have to fight apm/acpi to get it to sleep properly when I shut the lid, nor did I have to massage my wireless drivers and give them a little tune-up every week. It's a slick-looking block of aluminum, but it doesn't feel nearly as sturdy as the Thinkpad. I guess we will see if it holds up to my abuse and the many around-the-world trips this thing is going to have to endure. The battery status indicator on the battery itself is nice. The dock is cool-looking, but it quickly gets on your nerves. I found QuickSilver (<http://quicksilver.blacktree.com/>) to be a much quicker way to launch things without having to deal with the damn trackpad. Related to the last point, the very non-Unixy approach to installing applications simply by dragging them to the /Applications directory is very handy, especially the way everything is self-contained, but it does make it hard to run things from the command line. I tried the BBEdit demo, for example, and while I finally did figure out how to launch it from the shell, it wasn't very convenient. Not sure how you are supposed to use an editor that you can't quickly launch from the shell. I guess that's why QuickSilver exists. Vi/vim has worked fine for the past 20 years for me, and it will work fine for the next 20. Fink (<http://fink.sourceforge.net/>) makes me feel somewhat at home giving me apt-get for most of the common packages I use. It was trivial getting my development environment set up with Xcode and all the various libraries and tools I need to build PHP. You also quickly become familiar with Versiontracker (<http://www.versiontracker.com/macosx/>) for finding OSX applications. I like the fact that my ancient RG-1000 wireless gateways that I have a bunch of lying around in the garbage pile seem a lot more useful now since I can flash them with the Apple Airport firmware and make them look like real Airport AP's. It was always a massive pain trying to configure these things from Linux via the Java configurator thing. Expose is just cool!

The backlit keyboard has a high coolness factor although it seems a bit too dim to really be useful. It also doesn't light up the row of function keys. A way to turn it on on demand as well as a way to crank up the brightness would be nice. It may be hiding in there somewhere, but I haven't run across it yet. The light around the power plug is cool too, except this one, unlike the keyboard backlight is really bright and doesn't turn off when the machine goes turns off the lcd so twice now my wife has repositioned the idle laptop sitting in the bedroom out of sight. The Skype (<http://www.skype.com>) client is pretty and works very well. There is a plethora of decent web browsers. It came with Safari and IE. Add Mozilla, Firefox and Camino to that and you have 5 browsers that all

seem to work well. I am used to Firefox from Linux, so I am sticking with that.

The Bad (with some workarounds)

For some reason there is no way to do multiple virtual desktops in the standard GUI. Desktop (<http://wsmanager.sourceforge.net/>) manager solves that annoyance. Almost anyway. It would be much more useful if Apple-Tab only cycled through the applications on the current desktop and not all of them. I am guessing they can't hook in and change this without being more integrated with the GUI. (boxes on the left in the menubar image above)

iTunes is pretty nice, but why in the world doesn't it have a way to minimize it to the menu bar. If they are going to force me to keep that stupid menu bar on my screen at all times, the least they could do is make use of it. Luckily there is a very cheap little tool called Synergy (<http://wincent.com/a/products/synergy-classic/>) that fixes that problem. With Synergy installed using the "tabbed" (arrows in the menubar image above) look in the menu bar iTunes is great and would be in the "Good" section if I didn't have to buy this add-on Synergy thing to make it usable. It almost makes me want to pick up an ipod just to play with the

integration.

I was also surprised how badly it reacted when I fed it an xvid AVI file. I thought this thing was a multimedia monster. VLC (<http://www.videolan.org/vlc/download-macosx.html>) took care of that.

I hate IM with a passion, but unfortunately I need to use it. Fire (<http://fire.sourceforge.net/>) seems to do the trick, but I really haven't looked around much for anything better. My only real comment on this one is that the icon sure is ugly and that it hasn't crashed on me yet. Then again, it doesn't really understand Yahoo's status message stuff and Yahoo Messenger (<http://messenger.yahoo.com>) isn't too bad so I have been switching between these two.

For irc I am used to using XChat. The Aqua version (<http://xchataqua.sourceforge.net/>) isn't great, but it works. It really could use an update to the current xchat code. The tabs are centered, which is odd, and when you get disconnected from an SSL'ed connection it gets an error trying to reconnect. I have to restart it in order to get it to connect again. I should probably just install X11 and run the real xchat instead, but I haven't gotten around to that yet.

Pine from fink works nicely. But I have been trying to join this century by using a graphical email client. I don't think I will be able to though. Mail.app is really slow at dealing with huge mailboxes over IMAP. Thunderbird is better and I am close to being able to use it, but why doesn't it let me map 'R' to Reply-All? Using (<http://mozilla.dorando.at/keyconfig.xpi>) I can map it to Ctrl-R or Apple-R, but not simply R. And yes, I know this doesn't really have anything to do with the Powerbook or OSX, but since I am whining about stuff I figured I would throw it in. I also tried offlineimap (<http://offlineimap.sourceforge.net>) to try to speed up the IMAP and give me better disconnected support, but it uses Maildir and having 100,000 files in a directory apparently isn't something the OSX filesystem handles very well.

I have been using Kismet for years on Linux. Kismac (<http://binaervarianz.de/projekte/programmieren/kismac/>) is the OSX version and it seems a bit flaky. It has hung on me a few times and it needs to do weird driver swaps because the native OSX drivers don't support promiscuous mode. The built-in wireless card doesn't seem to do promiscuous mode at all even with the driver swap, but my Cisco and Orinoco pcmcia cards both work ok with it.

Apple left/right to switch between terminal windows is a nice touch, especially since Apple-Tab only cycles through a single window of each running application. But Terminal seems to be the only application that supports this. The same thing for Firefox would be very handy.

The Ugly

No dedicated Page-up/down keys! I never realized how much I used those until I didn't have them anymore. Fn-up/down is the same thing, but you need two hands for that.

Too many modifiers! Was Fn, Ctrl and Alt really not enough? Why an Apple key where the Alt key should be?

What's with the two Enter keys? Wouldn't it be more useful to have two Ctrl or two Alt keys instead? I definitely don't need two Apple and two Enter keys right next to each other. Obviously it is there for the 3 people left in the world that actually uses numlock and the keyboard as a numeric keypad replacement. So I have a proposal. Let's just deprecate numlock completely and give me a useful key there instead!

Inconsistency in keyboard access to various widgets. Some dropdown boxes don't have keyboard accelerators. 'u' to get you to "United States" in a long country dropdown, for example.

Not being able to hide the top menu bar is really hard to get used to. Coming from a 1600x1200 screen on the T42p down to this 1280x854 screen I am already feeling quite cramped, especially vertically so losing another 16 pixels to a mostly useless menu bar is annoying. You can't even change the font in it or do anything to make it smaller as far as I can tell. Would it be so bad to provide a way to autohide it and move/resize it? Or even better, let me dock it.

The damn clock widget in the menu bar won't show me the day of the month. You can toggle showing the day of the week along with AM/PM and 12/24 hour displays, but you can't get it to say "Dec.20 11:00". I am usually with it enough to know what day of the week it is, but I am always forgetting the day of the month. I know it shows up in the menu when you click on it, but that means I have to move the mouse, click and then hit escape instead of just glancing up there. There are replacement things like wclock that will do this of course, but having to run yet another process just for that minor thing is dumb.

To keep complaining about the clock widget, why doesn't it show something

useful, like the damn date, when I hover over it? Some applications, like Syngery, shows you something useful on hover from the menu bar, so I know it is technically possible. There are many other places where throwing in hover support would be nice.

I of course knew about the trackpad and single mouse button issue and I knew I would have problems with that, and I do. I was however under the impression that the GUI was designed in such a way that you really didn't need more than a single mouse button, but it turns out there are plenty of places where you really need to right-click (Ctrl-Click) on stuff. Like if you want to empty the trash without multiple clicks. A USB mouse mostly solves this annoyance, but that restricts me to having a surface nearby to use the mouse on.

I had heard iPhoto was slow. And it really is slow. It's slower than slow. I find myself pondering what it could be doing while it is chewing on my photos. I hope it is doing something worthwhile with my cpu, like curing cancer, while it is grinding away. It is good for organizing photos as long as you aren't in a hurry and with the iPhotoToGallery (<http://zwily.com/iphoto/index.xsl>) plugin it is really easy to keep the online photo album in synch as well. Sticking with the iPhoto gripes. When I import photos from a fast Sandisk Ultra 512M CF card in a PCMCIA adapter iTunes skips. This is a 1.5GHz CPU with a Gig of RAM. Can it really not copy a file from a CF card and play an MP3 at the same time? Given all my other gripes, this is probably the most disappointing.

Scrolling text input boxes in Firefox leave cursor artifact garbage on the screen. Probably something the Firefox folks are doing wrong and not Apple's fault. But still an annoyance I didn't have under Linux.

Why only a slow DVD-R drive? Every modern DVD burner out there these days doesn't care if you feed it DVD+R or DVD-R media. Seems like it is time to update that. And get one that isn't so noisy.

No amplified audio line-in. Means you need an amplified mic and all the cheap headsets out there designed for voice chat won't work. I suppose the answer is to use a USB headset instead.

Saturday, December 4, 2004

Philips Streamium SL300i

I seem to have this magic \$200 pricepoint barrier below which I buy anything that I can give an IP address on my LAN. A \$50 Philips discount plus a corporate partnership discount brought it under \$200 for me. It bugged me that my Linksys Media Appliance box only does Audio and Images and even though it has network support it doesn't make use of that to talk to the Internet in any way. Philips has thought things through a bit better. The Streamium can be used in 2 modes. They call it Internet and PC-Link. In PC-Link mode it is similar to the Linksys in that it uses Intel's UPnP Media Server protocol to fetch content from a local server and in Internet mode it connects to a number of online content services including Yahoo, Radio Free Virgin, Playhouse Radio, Andante and many others. You can also feed it your own streaming sources. Note though that it doesn't support WMA nor ASF so it can't stream from a number of Windows-centric streaming services.

It has both a 10/100 port and 802.11g support. I already have a WRT54G sitting on top of my TV so I just plugged it into that, but I did try the wireless config and it worked easily. Plugged in my WEP key and off it went. The Philips Media Manager software that comes with it is however horrible. It is this big Java monstrosity that chews up a whack of RAM and as far as I can tell only runs on Windows. Go figure. Luckily Philips didn't restrict the Streamium to their own media manager. You can install any UPnP media server you want. I prefer a little tiny one called TwonkyVision. It doesn't have a GUI but instead takes the sane approach and has a web interface. And it isn't written in Java so it is portable(!?) It runs well on OSX, Linux and Windows.

The TV interface seems a little clunky to me. It is somewhat hard to navigate and it seems to flicker a bit. I don't think it is me or the TV because the TiVo onscreen interface is rock-solid without any flicker at all. The online part of this at my.philips.com is also not great. It's not bad, but it seems like they could do a lot more with it. That's where you go to configure what the Internet portion of the device should connect to. I like the fact that no local machine needs to be available to use or configure the Internet side of it. Just go to the web site and choose which streams, video and photo services to use. Hook this thing up to an LCD hanging on a wall and configure it to show your Yahoo Photos or galleries from Lex Fletcher's Born to Shoot site and you have a pretty nice dynamic picture frame.

Prodding it just a little bit shows the following:

```
PORT STATE SERVICE
80/tcp open http
1720/tcp filtered H.323/Q.931
8080/tcp open http-proxy
```

And the web server returns:

```
200 OK
Cache-Control: no-cache
Connection: close
Server: Allegro-Software-RomPager/4.30
Content-Type: text/xml
Expires: Thu, 26 Oct 1995 00:00:00 GMT
Client-Date: Sat, 04 Dec 2004 16:11:14 GMT
Client-Peer: 192.168.1.106:80
Client-Response-Num: 1
```

followed by a bunch of XML which looks like a normal UPnP Media response. Interestingly enough it looks like other devices can use this one as their UPnP Media source. I am sure that is in the spec somewhere, but I never realized that

before that these things could be used as media relays.

All in all it is a cool little device.

Posted by Rasmus in Audio/Video at 07:05

Wednesday, November 24, 2004

Skype - Talk to your Laptop

Skype has been out for a couple of months now, but I only recently had a look. I have never been very impressed with the audio quality of these various voice-chat systems. However, Skype is a whole different story. With similar-looking clients for Windows, Linux, OSX and PocketPC, a mechanism called SkypeOut for making calls to regular phones and some nifty P2P principles applied to the problem of getting in and out from behind firewalls and NAT gateways all combining to create something that works extremely well for most people. The biggest problem people generally have with it is figuring out how to get a microphone installed and configured for their systems.

I installed it on a machine at home before heading to Paris and have been using it to talk to Christine from Paris. I also got my parents to install it by themselves and have had 3-way conference calls with me in Paris, my parents in Toronto and Christine in California. I have also called landlines in the US from Paris and Christine called my hotel landline in Paris using Skype. Overall the quality of all these calls were amazing. Every now and then there would be a slight drop or when 2 or 3 people all spoke at once it would occasionally garble things, but it is definitely the coolest piece of software I have used in a while. Although I may have to invest in a headset for it. You get some strange looks when you sit there talking to your laptop.

Posted by Rasmus in Software at 02:03

Friday, November 19. 2004

-1 T42p toy

A negative toy posting...

My nice new T42p was stolen by some loser at a PHP conference in Paris. It is amazingly inconvenient to lose a laptop like this. It was from inside the conference hall and there was virtually no non-geek traffic there. If a fellow geek actually stole my laptop from a PHP conference then there is something seriously wrong with the world. You can steal my car, my money, my shoes, I don't really care, but don't steal my damn laptop!

Now to determine what to replace it with. It took a while to get it configured nicely and I really don't have the time nor the energy to do that again. Perhaps a G4 Powerbook so I don't have to fiddle as much. Will be hard to give up that nice 1600x1200 Flexview display though.

Posted by Rasmus in WIFI Toys at 23:54

Saturday, September 11, 2004

Gallery and the Coral Distribution Network

The Coral Distribution Network (CDN) is a very nice shiny toy for all of us who sometimes struggle with limited bandwidth. Let the NSF pay for it!

My photo album is what chews up the most bandwidth from my site, so it made sense to rig it up first to optionally be available via CDN. CDN is basically just a big Akamai-like network of servers that cache things and serve them up to you from servers that are close to you network-wise. It has some fancy code behind it, but who cares how it works, it just does. Let's get to the interesting part. As you have probably figured out by now, you can access any site on the Web via Coral by simply appending ".nyud.net:8090" to the domain part of the URL and there are plugins available that add a menu item to automatically do this for you client-side. The problem with this is that you don't want to have to do this for every page on a site. It would be much nicer if the site would adjust its links such that if you enter the site via Coral, all the local links from that site would be Coral links. You may of course want to make some exceptions for pages that are interactive in some way, but for something like a photo album where 99% of people just look at mostly static content it works well.

You can see the patch at <http://gallery.menalto.com>

Note again that this quick hack does not in any way handle links that shouldn't be Coral'ed. Like the login link, comment stuff or all the administrative tools. It turns your Gallery into a read-only site if you come in using Coral. As the administrator of my album I can figure out that I need to go directly to it in order to add photos. But a more complete patch that doesn't Coralize stuff that shouldn't be would be nice.

You can have a look at it in action at <http://www.phpics.com.nyud.net:8090/>

I submitted it to the Gallery Folks so let's see if they take it and run with it to integrate it completely.

Update: A variation of this that also supports people running their Gallery on weird ports has been committed to Gallery's CVS, so you may just want to update from there to get it.

Posted by Rasmus in Software at 01:26

Thursday, May 20, 2004

IBM Thinkpad T42p

This should be a fun toy when it shows up. I ordered a T42p today with the 15" 1600x1200 screen and a Dothan 1.8GHz CPU. Kept the RAM and HD low and will just add more later. Crucial doesn't list it yet, but I figure it uses DDR PC2700 non-parity RAM just like the T41p. Now to find a decent 80G 7200 rpm notebook drive to toss in it.

I did consider a Powerbook for a while, but I am too used to Linux and I really am more comfortable on a Thinkpad. One of the big draws of the Powerbook was the DVI-out on it, but with the Thinkpad mini-dock which is only \$89, you get that anyway, and it's not like I will be tossing my 20" LCD into my backpack and bringing it with me, so the dock can just live permanently by the LCD for a very nice dual-headed workstation when I am working from home.

By the way, ordering this thing was amazingly painful. The ordering web site is/was completely messed up and there was no way to click your way to it. Had to call and have a human do it for me. If you are looking for one, I would suggest finding someone who works at IBM who will let you use their friends+family EPP discount and go in via www.ibm.com/shop/epp or if you own some IBM stock you can use the Shareholder Purchase Program at www.ibm.com/shop/us/spp to get a couple of hundred dollars off your price.

June 23rd Update: It finally arrived! Of course I am out of town so I can't play with it yet. Frustrating.

July 5th Update: Finally back in the country and have started playing with this beast. It's the same thickness as my old T20, about an inch wider and a bit over half an inch taller. But that Flexview display is amazing. And no, 1600x1200 looks just fine on it. I never really understood the argument that a display could be too small for a high resolution. Just set your font size to your liking. The higher resolution means your anti-aliased fonts have that much more definition to them making them clearer and easier to read which is exactly what you need on a "small" display.

I am waiting on another 512M of RAM and a speedy 7200 RPM 7K60 drive to install Debian on. I'll keep the 5K80 that came with it as a secondary XP drive that I can pop in the Ultrabay the one or two times a year I actually use Windows.

July 10 Update: Debian has gone onto this thing. It installed pretty smoothly. I always use this 31M XFS boot iso for installing Debian these days. To do a network install just remember to specify "e1000" when you get to the part that asks you which extra drivers to load. Here is my .config in case you are curious. I ended up using ATI's drivers for the FireGL T2 (basically a Radeon 9600 card) that is in it. There are also open source drivers ([here](#)) which work nicely, but the 3D acceleration wasn't very good. If you follow these excellent instructions it is easy to get the ATI drivers going, and you will have very fast accelerated 3D. Make sure you build your own modules instead of trying the precompiled ones he lists. I had a world of problems with those, but as soon as I built my own against my 2.6.7 kernel everything started working. I used the "fglrxconfig" program to generate my XF86Config-4 file for just single-headed 1600x1200 for now. Need to play more with the port replicator and dual-headed stuff and also come up with a way to reliably connect to 1024x768 projectors. I am getting around 1785FPS from glxgears and 380FPS from fgl_glxgears (default window sizes). Offscreen fgl_glxgears runs at 1125FPS.

Sound works fine with the snd_intel8x0 driver, and the built-in a/b/g wireless works nicely with the madwifi driver. I use apt-get to grab it via this entry in my /etc/apt/sources.list file:

```
deb-src ftp://debian.marlow.dk/ sid madwifi
```

Someone in the comments mentioned problems with pcmcia stuff, but I haven't seen any issues. Anything I plug in comes up right away.

Dual-booting to WinXP sitting on the original drive in the ultrabay worked on the first try. I added this to my /boot/grub/menu.lst file:

```
title Windows XP
map (hd0) (hd1)
map (hd1) (hd0)
rootnoverify (hd1,0)
chainloader +1
```

The big thing I still need to work more on is ACPI and getting it to suspend and wake back up. It suspends perfectly right now, but it just won't come back out of suspend which makes the fact that it can suspend much less interesting. Another interesting problem I hit was that if I used the Radeon Framebuffer to get a cool-looking console then the fglrx ATI driver would crash the system on switching between X and the console. If I don't use the framebuffer for the console everything is fine. Haven't tracked down a solution to this one. For now I just use the vesa framebuffer for the console which works well.

July 17 Update: I spend half my life on planes and the other half presenting. I haven't found any way to make the former easier on me as I absolutely hate flying, but for the latter I trawled the Net and came up with an idea by Klaus Weidner for running a vncserver and then a viewer onto that server session both on the local lcd and on the external vga port. That means that now when I present I can have the contents of the projector in a window on my desktop. This will be very nice, especially for my duller talks as I can read email or irc while presenting without people seeing that.

First, here is my XF86Config-4 file. Note the dual fglrx device sections and the dual screen sections and finally the single and dual ServerLayout sections. Unfortunately X is quite unhappy starting up with the dual layout if nothing is connected, but you can check that with a tpctl call. I use this little startx wrapper script:

```
#!/bin/sh
if [ `tpctl --id | grep "monitor type" | cut -c41` != 0 ]; then
    startx -- -layout dual;
else
    startx -- -layout single;
fi
```

So I just need to restart X to have it automatically figure out if the second display should be enabled or not. Next, to run vncserver and the viewers along with a window manager (metacity) and a panel I use this script:

```
#!/bin/sh
PWFILe=$HOME/.vnc/passwd

vncserver -geometry 1024x768 :3
sleep 1
xvncviewer -passwd $PWFILe -shared -fullscreen -display :0.1 :3 &
x2vnc -passwdfile $PWFILe -shared -east localhost:3 &
xvncviewer -passwd $PWFILe -shared :3 &
DISPLAY=:3
metacity &
gnome-panel &
```

As far as my suspend problems go. The problem is the ATI fglrx driver. I would have to switch back to the radeon driver but then I would lose tv-out and some 3d-performance. Probably not a bad tradeoff actually.

Posted by Rasmus in WIFI Toys at 20:07

Sunday, March 14, 2004

Buying a new DVD player

Carl managed to break our old Sony DVD player which of course meant I got to buy a new toy. I was very tempted by the Gateway Networked DVD Player and the new LiteOn LVD-2010 which is also a networked player. Being able to stream things to the player skipping the middle step of having to burn to DVD first is attractive, and of course, why would you buy any powered device without a NIC? Every device should have an IP, no matter how inappropriate. However, I played a bit with the Gateway player and the menus seemed clunky and the number of formats it could play natively was limited. If the machine you are streaming from has enough juice you can transcode from a large number of formats on the fly, but for some reason I still wasn't sold on it. And the LiteOn doesn't seem readily available yet. I couldn't find it anywhere. So instead I picked up a Philips DVD727 for \$79 from Fry's. It is a slick little progressive scan player (not that I have a TV capable of that yet) with S-Video, component and optical outputs on the back. The menus are simple and making it region-free was trivial. Open tray, 9 9 9 0 (slowly), close tray. I tested it with a Danish region 2 PAL DVD and it played perfectly. Unlike many other players you can make region-free, this one you can set to any region, so if you have any region-locked dvds that won't play in region 0 you can set it explicitly to the appropriate region. The next test was to see how well it handled a DVD+R data disc. I grabbed some 3500 photos of Carl, 30 videos of various AVI, WMV and MPEG formats and a couple of hundred mp3s and stuck them in Photos/, Videos/ and Music/ directories on a regular data dvd and popped it into the player. It came up with a nice little directory of what was on the DVD. It was only able to show the first 650 jpegs in my Photos directory. The manual says there is a limit of 500 files per directory. So it looks like I will have to spread my 3500 pictures out over many sub-directories. Annoying, but not a big deal. The mp3 playback is nice. It has playlists and even a shuffle mode. You can also start the mp3s playing and switch over to the jpeg slideshow and have them both going at the same time. For the videos it was only able to play the mpeg files. No AVI nor WMV support apparently. But it still leaves our previous Sony player in the dust by not needing to do full DVD authoring which takes forever. Being able to browse your backup DVDs in your DVD player is very nice. This player also supports SVCD and regular CD-R of course.

Posted by Rasmus in Audio/Video at 23:15

Friday, March 5, 2004

Kismet on the Linksys WRT54G

The little Linksys WRT54G box is a terrific generic Linux platform to run just about any networking code on. I have found that the radio on it when cranked up to its full 84mw is better than any of my pcmcia cards including the 100mw Cisco-350 I normally use when I need to pick up some distant signal. I have this 5dbi Maxrad antenna I normally use with the Cisco card and even with that it doesn't match the sensitivity of the WRT54G with the stock antennas. I also picked up a dual diversity flat-patch antenna from Hyperlink to see if I could extend my packet detection range a bit. There is a picture of it in the extended entry. Also note the updated section at the end of the 54G and no Wires post.

For those that haven't run across it before, Kismet is a very handy 802.11 monitoring program which is used to detect wireless activity.

There is a MIPS binary for kismet_drone and kismet_monitor at <http://gattaca.ru/~nikki/wrt54g/kismet.tar.bz2>.

To get it up and running, first you need command-line access to your gateway. I suggest sticking this firmware on it. Just unzip and use the standard "upgrade firmware" option to switch to it. Reboot the box and under the Administration menu turn on telnet and under the wireless menu put it into Client mode. Uncompress the kismet tarball on some machine, telnet into the gateway and from /tmp either scp or wget the files into /tmp/kismet/bin and /tmp/kismet/etc. Edit the /tmp/kismet/etc/kismet_drone.conf file and make sure you pick the right source ethernet device based on your wrt version. For version 1.0 and 1.1 use eth2 and for a v2 gateway, use eth1.

```
# WRT v1, v1.1
```

```
source=wrt54g,eth2,wrt54g
```

```
# WRT v2
```

```
#source=wrt54g,eth1,wrt54g
```

To run it, first make sure you are not associated with a gateway already. It will actually still work, but it won't channel hop automatically. Also a good idea to make sure you don't send out any probes by sticking it into passive mode. I would suggest these steps:

```
wl disassoc
```

```
wl passive
```

```
wl scan
```

```
wl scanresults
```

The scan and scanresults is just to get a sense of whether there is anything out there. It will tell you if it sees any gateways and what their signal strengths (rssi) are. Here is the typical output from one of my gateways:

```
# wl scan
```

```
# wl scanresults
```

```
SSID: "Canada"
```

```
Mode: Managed RSSI: -40 dBm noise: -82 dBm Channel: 3
```

```
BSSID: 00:06:25:C5:32:21 Capability: ESS WEP ShortSlot
```

```
Supported Rates: [ 1(b) 2(b) 5.5(b) 11(b) 18 24 36 54 6 9 12 48 ]
```

```
SSID: "Canada"
```

```
Mode: Managed RSSI: -71 dBm noise: -82 dBm Channel: 3
```

```
BSSID: 00:0C:41:D3:99:E1 Capability: ESS WEP ShortSlot
```

```
Supported Rates: [ 1(b) 2(b) 5.5(b) 11(b) 18 24 36 54 6 9 12 48 ]
```

Now to run the drone, do this:

```
/tmp/kismet/bin/kismet_drone
```

You should see something like this:

```
Suid priv-dropping disabled. This may not be secure.
```

```
No specific sources given to be enabled, all will be enabled.
```

```
Enabling channel hopping.
```

```
Disabling channel splitting.
```

```
Source 0 (wrt54g): Enabling monitor mode for wrt54g source interface eth2 channel 6...
```

```
Source 0 (wrt54g): Opening wrt54g source interface eth2...
```

```
Kismet Drone 3.1.0 (Kismet)
```

```
Listening on port 3501 (protocol 8).
```

```
Allowing connections from 192.168.0.0/255.255.0.0
```

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

And now on your Linux box which should be connected directly to one of the switch ports on the gateway with an appropriate ip allocated for both gateway and Linux box, of course, find your kismet.conf file and put this in it:
source=kismet_drone,192.168.1.3:3501,drone

Now you are ready to fire up kismet. If everything worked and there are gateways out there you should see something like this:

Here you see my two other wrt gateways each with an essid of Canada, an mlife access point somewhere, one named WesClark(?!) and one named default. The colours indicate if they are using encryption and generally how secure they might be. Green means encryption is used, yellow means no encryption, but at least the default config has been changed in some way so it may not be trivial to access it and red means a gateway which is still running with its default wide-open config. Here is the WRT with the Hyperlink antenna pointing out a window.

Posted by Rasmus in WIFI Toys at 18:52

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Saturday, February 14, 2004

54G and No Wires!

My home network has 3 annoying problems:

#1 - The house has an odd shape and Christine's office is a long way from the spot the cable comes into the house and where the wireless gateway is. The signal strength in her office is barely adequate.

#2 - I have a nice Raid5 server which I used to keep in the garage at the previous house, but in this one the Internet connection terminates inside the house and the machine is too noisy and ugly to have sitting in the dining room. I'd like to get it back into the garage, but there is no connection there currently.

#3 - There is no phone jack anywhere near the TV, so I am periodically running a long phoneline to the Tivo so it can make its daily call. This is one of old series-1 Tivos which has one of the first TivoNet cards in it, so if I could get network connectivity to the TV I could make it update via that instead.

The solution is of course to buy more toys! In this case a couple more WRT54G wireless gateways. At \$79.99 from Amazon they are relatively cheap and Linksys has released all the source to the Linux-based firmware. You can get it from their GPL Code Center. People have taken this code and created their own customized firmware. The best right now is from SveaSoft.

Here is a picture of what I'd like to build:

The main AP (192.168.1.1) is the one that is there now. The secondary (192.168.1.2) is the one in the garage that I will plug my RAID5 box into solving problem #2. Christine's office is also above the garage, so this secondary AP will be the one she will associate with and hopefully get a stronger signal. That means it will have to talk to the main AP via WDS (Wireless Distribution System) and then turn around and talk to any wireless clients. The radio will have to flip back and forth and I unfortunately lose half the bandwidth that way. This is 802.11g though and currently trying to go directly from that office to the main AP on the weak signal is causing the connection to drop back to 1Mbps as it is. I am hoping to see close to 10Mbps in this new setup from this AP in repeater mode.

To solve problem #3 I am going to have a third wrt54g near the TV. This one doesn't need to repeat since it isn't that far from the main AP. All I need from it is to connect as a regular wireless client to the main AP and then act as a switch where I can plug the Tivo and probably the WMA into. The WMA has wireless, but it is only 802.11b and this way I can get more bandwidth to it by connection it to this 802.11g connected switch.

At this point I have AP #2 installed in the garage and it is working well. I ran some 30-second iperf tests on it:

AP2 = WRT54G v2.0 Satori-pre1 (AP mode w/ 40-bit WEP)

AP1 = WRT54G v1.0 Satori-pre1 (AP mode w/ 40-bit WEP)

LAN = Various Linux servers and an XP box all with 100M NICs

WAN = Thinkpad T20, Linux 2.6.3-rc2, Netgear WG511 802.11g card with the prism54 driver.

LAN-AP2-WDS-AP1-LAN	9.2 Mbits/sec
LAN-AP1-LAN	93.8 Mbits/sec
LAN-AP2-LAN	93.9 Mbits/sec
WAN-AP2-LAN	19.5 Mbits/sec
WAN-AP1-LAN	19.9 Mbits/sec
WAN-AP2-WDS-AP1-LAN	5.1 Mbits/sec
WAN-AP1-WDS-AP2-LAN	5.8 Mbits/sec

A note on the above performance numbers. There are quite a few walls between AP1 and AP2 for these measurements, so the WDS speeds are not what they could be. I tested them next to each other as well and was able to get it up to about 14 Mbits/sec. By locking it down to g-only, turning off WEP and fixing the speed to 36Mb/s I was able to get it up to 17 Mbits/sec. It's still not as fast as I would like and I think I may just get a better antenna for the main AP and have Christine connect directly to that while running the AP in the garage in client mode. This should be quite a bit faster.

I have ordered another WRT54G to sit by the TV, although I am also tempted to get a decent antenna and have one sit around in monitor mode just to keep track of what other traffic is flowing by the house here.

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

This Dual Diversity Flat Patch Antenna from HyperLink looks like it would be suitable to get a stronger signal to the far end of the house. Remember that the Linksys boxes come with this weird RP-TNC connector, so any antenna you buy for it should have a Female Reverse Polarity TNC on it.

[Update March 5, 2004] I picked up a couple of those Hyperlink antennas and they didn't boost things as much as I had expected. I did hook up the WRT behind the TV as per the network diagram above and am running it in client mode and it works well. It does seem like there is an issue with plugging multiple clients into it in client mode, so probably a driver bug somewhere. Nothing really to getting client mode working. Simply set it to client mode on the wireless tab of the firmware. However, I find that for some reason mine isn't associating automatically. Logging in and manually doing:

```
wl join Canada key 3132333435
```

seems to do the trick. (Not my real key obviously) You can then check the status with:

```
# wl assoc
```

```
SSID: "Canada"
```

```
Mode: Managed RSSI: -40 dBm noise: -85 dBm Channel: 3
```

```
BSSID: 00:06:25:C5:32:21 Capability: ESS WEP ShortSlot
```

```
Supported Rates: [ 1(b) 2(b) 5.5(b) 6 9 11(b) 12 18 24 36 48 54 ]
```

One very useful feature of client mode would be to bring it on trips. With its superior radio combined with one of the Hyperlink antennas it should be able to pick up open gateways at quite a distance. See the Kismet post for further details on using it for finding gateways and with client mode you could then associate the wrt with a gateway and plug yourself into one of the wired ports. No more messing around with wireless drivers on your laptop.

WDS mode between AP1 and AP2 is all done through the web interface as well. Each WDS endpoint needs an ip. 10.0.0.1 and 10.0.0.2 in my case. Then specify the MAC of the other AP in each web interface. Make sure both are set to the same channel and same ESSID and turn off any firewalling. There is one step the current firmware doesn't handle, so you need to do that manually after each reboot. Either via the Administration->Diagnostics tab in the web interface or by logging in, issue this:

```
brctl addif br0 wds0.2
```

to add the wds link to the LAN bridge. After doing these steps you should be able to ping the other side.

Posted by Rasmus in WIFI Toys at 02:35

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Monday, January 12. 2004

Spamassassin

...is normally a very good toy. I use it extensively with Razor2 and Bayes hooked into it, it catches nearly everything. Today, however, I was suddenly inundated with Viagra spam. Odd, since any such spam usually sets off all sorts of SA alarms. And sure enough, the alarms were there:

```
X-Spam-Status: No, hits=0.9 required=5.0 tests=BAYES_99,BIZ_TLD,CLICK_BELOW,
HABEAS_SWE,HTML_50_60,HTML_LINK_CLICK_HERE,HTML_MESSAGE,
MIME_HTML_ONLY,MIME_HTML_ONLY_MULTI
```

So what gives? I have my BAYES_99 rule cranked way up to 5.4, so that alone should have put it over the 5.0 spam requirement, never mind all the other ones it triggered. The answer is of course that there is a large negative rule in there. In this case it was this HABEAS_SWE thing. The default SpamAssAss scores file has:

```
score HABEAS_SWE -8.0
score HABEAS_VIOLATOR 16.0
```

And it turns out that Habeas is some sort of "good spam" company that you can pay to get yourself whitelisted if you really need to spam people. I can see how that could be useful if you have a newsletter or something that people subscribe to and then it can't get through because of filters, but then these Habeas people damn well better be on the ball and triple-check the intentions of everyone and also run a tight ship security-wise. Given the fact that I received at least 20 Viagra spams before I killed that -8 rule, they obviously weren't quite on the ball and I don't particularly appreciate that this rule was in the default SA config to begin with. I haven't tracked down exactly who put that rule in and what sort of compensation changed hands. If someone knows, I'd like to know.

I would suggest that you find your local.cf file. Mine is in /etc/spamassassin/local.cf and add:

```
score HABEAS_SWE 0.0
```

Making it neutral.

Posted by Rasmus in Software at 03:47

Tuesday, December 16, 2003

Linksys Wireless Media Appliance [updated]

[update]

As Sean Lincolne pointed out, squishguava, the boot image this thing loads over the network when it boots, is just a simple cramfs image. The "COMPRESSED ROMFS" in my earlier 'strings' is a dead give-away on this. I finally got some time to have a peek at it. It has the following in it:

```
libc-2.1.3.so*libm-2.1.3.so*librms.so*rio*libdl-2.1.3.so*libmp3msp.so*libthreadutil.so*scripts/libhttpmsp.so*libmrdi.so*libw  
mamsp.so*version.txtlibiupnp.so*liboalmsp.so*mrd*web/libixml.so*libpil.so*mrdDevice*liblmsp.so*libpthread-0.8.so*pirre  
g.dat*
```

It is obviously not a full Linux filesystem, as can be seen by looking at files in the scripts directory, this filesystem gets mounted on /guava. I compiled a small statically linked busybox and stuck it in bb/ which among other things had bb/sbin/telnetd in it. Then I hacked up scripts/rio-script and had it launch the telnetd right before it launches the rio. Built a new cramfs squishguava image and booted with it. It actually booted just fine. I was worried they might have some sort of checksum check, or my added 600k would overflow something, but it was fine. The bad news is that I couldn't get in via telnet. So, either, my telnetd didn't start properly, or the port is blocked.

I did notice that the default shell is /bin/ash which also happens to be the default shell in BusyBox and Linksys has used Busybox on other devices in the past. I bet the rom firmware which has the root filesystem for this thing is just a busybox image. So a bit more hacking on that rio-script should let me either somehow get a message out to me by trying various standard busybox commands, or I can run some stuff to try to deblock the port. Any suggestions on what is likely to work right away?

More info. I replaced the rio binary with my arm cross-compiled telnetd binary and it then doesn't get beyond the "Launching remote-IO" message during boot. At least it tells me that what I am doing has some effect. But I still can't get in via telnet. I also tried replacing it with a script that tried to ping out, cat stuff to /dev/dsp and echo stuff to various devices and none of that did anything I could see/hear.

[/update]

This looks like a nifty little box that will make it easy to access mp3's and photos directly from a remote-control TV-displayed interface. Much nicer than needing to stick a PC next to the TV/Stereo in the living room.

This little device showed up today. Had no trouble configuring it and hooking it up once I shuffled the various cables around a bit on the back of the TV and stereo. The music navigator is really nice on it and I like that you can play mp3's while a photo album is cycling through. Will have to try this thing against some Samba shares later on.

No luck on the Samba shares, or any sort of shares at all actually. I did a bit of sniffing of the datastream between the WMA and XP. When it starts up the first thing it does after getting an IP via DHCP is to grab its OS image from the XP box. That image is clearly a Linux 2.4.17 kernel and all communications appear to be via a UPnP A/V Media Renderer SOAP thing. As far as I can tell, when you designate a directory via the Media tool on the XP box, it creates a regular .m3u playlist out of that and serves it up to the WMA when requested. There doesn't appear to be any encryption involved, so getting this thing to work with a Linux box as the server would involve creating a UPnP SOAP server that understood the requests from the WMA. Not that this is a trivial effort, but certainly not impossible and once done this thing would be able to serve files up from anywhere a Linux box could access files from. Frankly I don't see why the heck the SOAP server they provide for XP can't serve up its playlists from a network share. There doesn't appear to be any technical reason for this restriction. I bet that with a bit of hacking and with the help of libupnp this is quite feasible.

Or, alternatively, create a custom image from the sources Linksys is supposed to provide. They have their GPL Page but it doesn't list the WMA11B (yet?). As George notes, SOAP isn't exactly ideal for something as simple as moving mp3s and image files around. An alternate image that was able to mount shares directly, would be cool. It might require sticking a .m3u playlist file in each directory so you wouldn't need to do that on the WMA, but that wouldn't bug me either.

For more info on the technology in this device, have a read through Intel's Digital Home Site or see the extended entry for some nitty-gritty protocol details.

Don't ask me why the boot image is called squishguava, but it is. Can't gleam too much out of it other than the fact that it is very likely to be a Linux image.

% strings squishguava
Compressed ROMFS
6i8V

Compressed
version.txt
libhttpmso.so
libiupnp.so
libixml.so
libllmso.so
libmp3mso.so
libmrdi.so
liboalmso.so
libpil.so
librms.so
libthreadutil.so
libwmamso.so
mrdDevice
pilreg.dat
libc-2.1.3.so
libdl-2.1.3.so
libm-2.1.3.so
libpthread-0.8.so
scripts
channelmgrscpd.xml
remoteinputscpd.xml
riodevicedesc.xml
rioscpd.xml
ConnectionManager.xml
MediaRendererDevDesc.xml
RendererControl.xml
TCxml.xml
Transport.xml
rio-script
mrd-script

The data stream when you fire this thing up is much more telling. Near the start we see a, "Hey I am Awake" message. 172.16.10.100 is the WinXP box and 172.16.10.105 is the WMA.

```
NOTIFY /AdapterInfoService/event HTTP/1.1
HOST: 172.16.10.100:8037
Content-Type: text/xml
NT: upnp:event
NTS: upnp:propchange
SID: uuid:1
SEQ: 0
Content-Length: 324
<firmware>
  <version>Ver. 11 R06</version>
  <time>09:50:28 AM</time>
  <date>08/01/03</date>
</firmware>
```

It also sends a NOT_STARTED message:

```
NOTIFY /ApplicationTransferService/event HTTP/1.1
HOST: 172.16.10.100:8037
Content-Type: text/xml
NT: upnp:event
NTS: upnp:propchange
SID: uuid:2
SEQ: 0
```

Content-Length: 177

NOT_STARTED

And the response back from the WinXP box is:

POST /AdapterInfoService/control HTTP/1.1
HOST: 172.16.10.105:62063
SOAPACTION: "urn:schemas-upnp-org:service:AdapterInfoService:1#GetExtDeviceDescription"
CONTENT-TYPE: text/xml ; charset="utf-8"
Content-Length: 303

Posted by Rasmus in WIFI Toys at 23:25

XS-Drive

This XS-Drive toy looks useful. I noticed Compgeeks has them for \$89.99 without a drive. Then it is just a matter of picking up any 9.5mm laptop drive and slapping it in there. DPReview of it

On the other hand, it does seem a bit silly walking around with two harddrives all the time. One for your mp3 player and one for offloading your digital cameras. So they have a combined version called the XS Drive Pro. I wonder how well this one works. Very little info out there on this thing and it looks like it is only selling in the UK right now. If anybody has one, please report back.

Posted by Rasmus at 21:15

Sunday, November 23, 2003

Looking for an 802.11g bridge

I've given up on the idea of a PCI 802.11g card for the Linux server. Driver issues are too annoying and an 802.11g bridge isn't much more expensive anyway. So now the problem becomes finding a good bridge. The Netgear WGE 101 seems like it might do the trick. At around \$95 it is also one of the cheaper ones.

The WET54G is the Linksys bridge. Reviews are not very flattering and it is more expensive.

The WGA54G is the Linksys game adapter which doesn't have anything to do with games, of course. It just means it is a single-port bridge and it is a bit cheaper than the WET54G. But unless it was much cheaper than the Netgear I think I'd still prefer the Netgear.

Another option would be to pick up another WRT54G since it can be hacked into acting as a bridge if you turn on WDS support.

And finally there is the D-Link DWL-G810. I have not had good luck with D-Link gear in the past and it also looks rather ugly. Anybody out there have one of these?

Posted by Rasmus in WIFI Toys at 10:15

Tuesday, November 11, 2003

802.11g Netgear WG511 and Linux

I picked up a cheap Netgear WG511 the other day. Got it for \$35, probably because they have recently released the WG511T which uses the Atheros super-G chipset. The older WG511 uses the Prism Duetto chipset which isn't officially supported on Linux by anybody. I say officially, because some code has snuck out and there is a new site out there devoted to it. Have a look at <http://prism54.org/>. I haven't tried that driver yet, but I will update this when I do. [Update - Feb.18/2004] I am now using the driver from prism54.org compiled into my 2.6.3 kernel on my Thinkpad and it works nicely.

For now I wanted to give the Linuxant Driverloader a whirl to see if I could use the native Windows XP drivers directly on my Thinkpad with a very recent 2.4.22 kernel. It worked amazingly well. See the extended entry for the step-by-step screenshots.

Of course, the whole point of going with 802.11g over 802.11b is to go faster. I haven't done any real performance tests yet with this Windows driver running on Linux. Hopefully I will get some time to test it against the native driver soon. Step 1 was to plug my Thinkpad into a wired port. How old-fashioned! And then plug the new Netgear PCMCIA card in. My kernel obviously didn't know what to do with it at this point.

I then grabbed the driverloader-1.38.tar.gz tarball, ran "make install" and then the dldrconfig command as shown:

From then on it was a web-based install. Cool!

So the first hurdle was to find the Windows XP drivers for the WG511 and actually get the .inf, .sys and .arm files out of the annoying executable Netgear provides. I cheated and used an XP box to install them and just copied them over from the drivers directory. They are probably also on the CD that came with the card, but I wanted the latest. You then feed the web interface the .inf file.

It figures out that I need the .sys and .arm files as well.

It has ingested the Windows driver and reads the MAC off of my card.

Ah, an Advanced button. I like those. You always find all the essential settings that the vendors think you are too dumb to understand there.

Here we find that we can enable the power saving features of the driver.

Next I need a free trial license to activate it. Clicking through (remember I have a wired interface up still) is easy enough. Just enter the email address and license string you get from the Linuxant site:

And you are done!

Now just use your standard iwconfig tool like with any other wireless driver and it just works!

Posted by Rasmus in WIFI Toys at 23:10

Canon Digital Cameras

Finally a digital SLR we can afford. Called the Digital Rebel in the US and EOS 300D elsewhere. They left out a few things, but still a very impressive camera.

DPReview The S400 is still the coolest small 4MP digital camera around and a perfect sidekick to the Digital Rebel. We use it as our travel camera to go with the G3. The G3 with the external flash attached and with its better lens still takes noticeably better pictures though.

DPReview

Although the newer SD-memory SD100 is cool as well.

DPReview

And an even newer SD-based camera. The SD10 loses the optical viewfinder and the optical zoom, but it is really really small and comes in 4 colours.

DCResource Review

This is our current digital camera. Aside from being a bit bulky, it is a great camera.

DPReview

Posted by Rasmus in Digital Cameras at 06:11

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Monday, November 10, 2003

Tivo!

Sophie asked me about getting a Tivo. Basically if you don't do DirecTV, a series 1 Tivo is all you really need. I have had an original Phillips 14-hour series 1 Tivo for years. Stuck a 120G drive in it and a network card and it is a happy little box. Should be in some sort of hall of fame for the best geek gadgets of all time.

ComputerGeeks has the original Tivo already upgraded with an 80G drive for \$219 which isn't a bad price. And if you add the promotional code "GEEKTIVO" during checkout, you get \$20 off. \$199 for a series 1 80G Tivo is not bad.

The original 14-hour Phillips Tivo can sometimes be had for \$89 from ServiceDVR. But right now they don't seem to have any of the 14-hour ones. The cheapest is a 30-hour for \$129. So, for the geeks out there, find yourself a 14-hour series 1 Tivo in good working order and upgrade it. Plenty of people are upgrading to series 2 or DirecTivo boxes so you can also check Ebay and look for an HDR-112 (original 14-hour Tivo) or better.

For the non-geeks, that \$199 Tivo from ComputerGeeks is probably your best option right now. Also have a look at Seth's TiVo Evangelism Page for more Tivomania.

Posted by Rasmus in Storage at 19:27

Sunday, November 9, 2003

802.11g PCI card options

Buffalo WLI-PCI-G54 uses the Broadcom chipset and has a cool-looking external antenna. At this point I think the only hope of getting this to work with Linux would be through Linuxant's driverloader.

D-Link DWL-G520 uses the Atheros 5002 chipset. Should work with the madwifi driver, or with the Linuxant driverloader. It also claims to support 108 Mbps Extreme-G.

Netgear WG311 uses the Intersil Prism GT chipset which has no native Linux driver that I know of. But the Linuxant driverloader says it supports it.

Linksys WMP54G uses the Broadcom chipset and should work with the Linuxant driverloader.

The Linksys WMP55AG is an a/b/g card whereas all the previous were just b/g. This one uses the Atheros 5212a chipset and should work with both the madwifi driver and the Linuxant driverloader. This one is actually just a PCI card with a mini-PCI adapter on it with a mini-PCI card plugged into it.

D-Link DWL-AG520 uses an Atheros chipset and should be supported by both the madwifi and the Linuxant driverloader. Like the Buffalo, it has a nice beefy external antenna.

Netgear WAG311 like the other a/b/g cards is Atheros-based so it should work with both the madwifi and the Linuxant driverloader. It probably would be a good idea to get an Atheros Super A-G based board so I can go 108Mbps when the drivers support it and when I get a gateway that can go that fast. I don't think my Broadcom-based WRT54G is going to be able to support Super-G. I think the Linksys, D-Link and Netgear a/b/g cards are all based on the same Atheros chipset, so the only deciding factor is likely to be price between these. If anybody has one of these and can inject a bit more data it would be appreciated. I will update and bump this up as I learn more.

Posted by Rasmus in WIFI Toys at 22:02

Wireless Video Camera

Another Linksys gadget. A motion-sensitive wireless video camera that can alert you via email when it sees motion and it streams out 320x240 video. Could use it to bring CarlCam back, although my outbound bandwidth on this Comcast cable connection isn't really sufficient for it.

I have been eyeing the networked cameras like the Axis 2100 (on the right) for quite a while, but they always seemed too pricy. This Linksys is \$100 cheaper than the Axis and it is wireless.

Posted by Rasmus in WIFI Toys at 08:37

USB Geek Watch

There is something cool about being able to store 256M of stuff in your watch. I don't have one, and I doubt I would use it much if I did. I am not that keen on watches. Being somewhat useless just makes it more cool though.

Posted by Rasmus in Storage at 08:11

Blog Export: Rasmus' Toys Page, <http://toys.lerdorf.com/>

Saturday, November 8, 2003

Linksys WRT54G Router

This is my current wireless router. Not because I love all things Linksys, because really I don't, but because this is a neat little 125 MHz MIPS box running Linux. And because of an oversight by Linksys in the Ping tool on their Admin page it is easily hacked.

You can do things like running Snort directly on it and have your gateway email you or notify you via irc/IM/pager if someone is trying to sniff your network. You can of course also run fancy netfilter/iptables rules or anything else you can typically do on a Linux box when you have it acting as your gateway.

See the SeattleWireless page on the wrt54g for all the details.

Posted by Rasmus in WIFI Toys at 02:06

Cheap 802.11b gateway

If you look around a bit, you can find this Netgear MR814 wireless gateway for next to nothing. \$35 on Amazon right now, for example, but often even cheaper. I bought one a while ago but didn't use it because it was very buggy. However with the latest firmware it is actually pretty good now. But, also consider that 802.11g gateways are getting cheaper quickly, so I am not sure there is much point buying an 11b one anymore.

To upgrade it, make sure you clear your previous config completely by holding the little reset button on the back for 10-15 seconds, then install this firmware: MR814v1414RC3.zip

Posted by Rasmus in WIFI Toys at 02:04

Our current mp3 players

I like this little iRiver 256M flash player. Small and light and great for running. Not very hackable though. The latest firmware does let you manage it like normal removable flash storage so you can use it easily from Linux or on Windows without any special software.

TechTV Review

Firmware upgrade info

The Archos harddrive players are getting a bit dated, but that also means they are cheap. I picked up a 10G Studio player for \$85 a while back. It only has USB1.1 and a tiny 2-line display which makes it nearly impossible to navigate your songs, but I just stick it on random and let it do its thing. The best part about it is that it is very hackable. Check out the open source RockBox firmware for it:

[Update March 10, 2004] The Archos hard drive won't spin up anymore. Guess it took a few too many beatings coming with me on all my runs. Waiting to see if anything interesting comes out soon to replace it.

Posted by Rasmus in Audio/Video at 02:03

Cool PHP Apps

People are always asking me which PHP applications I prefer. It is obviously a very subjective thing, but here is a list of ones I have worked with and liked:

FUDForum is a very nice forum package writtern by Ilia.

Gallery is the ultimate photo album application. You can see it in action at phpics.com

Serendipity is a cool blog package. I really don't like blogs, but as you can see from this toy page which is using

Serendipity it doesn't have to be just for your standard boring and useless blog. It can be used for a useless toy page too.

Posted by Rasmus in Software at 02:01

Dell 2000FP 1600x1200 LCD Monitor

I have this Dell LCD and I highly recommend it. It's rather bulky and not nearly as slick looking as many others, but the image quality is great. Wait for a Dell Small Business sale and you can pick it up for \$750-\$850.

Make sure you get a DVI-capable graphics card to go with it. I found DVI quality quite a bit better than analog.

Review: Cadalyst Labs

Posted by Rasmus in Displays at 02:00

New Zaurus C760

I was in Japan recently and this little guy had everyone excited. With its swivel screen and full qwerty keyboard it is an interesting little Linux box. Although at \$800 currently I am not rushing out to buy one at this point.

InfoSyncWorld Review

Posted by Rasmus in Phones and PDA's at 01:43

Treo600

I am not a huge fan of PDA's and cellphones. But this new Treo is pretty cool. When it comes down in price a little bit I will probably pick one up. Being able to ssh and irc anywhere from a little device with a full keyboard is cool.

Posted by Rasmus in Phones and PDA's at 01:41

Canon Flash

A nice E-TTL compatible flash which works perfectly with our G3 and it will work nicely with the Digital Rebel when/if we pick that up. It can also act as a slave flash for the much more expensive 550EX.

I find my pictures are sharper and more alive when I am taking any sort of indoor or otherwise low-light shots with this flash.

Posted by Rasmus in Digital Cameras at 01:39