

**CI HOST**  
MANAGED WEB HOSTING

**\$59 a month!**

**Wow!**  
[click here](#)



- [faq](#)
- [code](#)
- [awards](#)
- [journals](#)
- [older stuff](#)
- [rob's page](#)
- [preferences](#)
- [submit story](#)
- [advertising](#)
- [supporters](#)
- [past polls](#)
- [topics](#)
- [about](#)
- [bugs](#)
- [jobs](#)
- [hof](#)

[Cringly Wants A Supercomputer in Every Garage](#) | [Login/Create an Account](#) | [Top](#) | [276 comments](#) | [Starting at #50](#) | [Search Discussion](#)

Threshold:   Nested  Oldest First

The Fine PrintThe following comments are owned by whoever posted them. We are not responsible for them in any way.

[1](#) | (2) (Slashdot Overload: CommentLimit 50)

**My Hank Dietz (creator of KLAT2) story** (Score:4, Interesting)  
by IvyMike on Thursday December 27, @01:33AM (#2754188)  
([User #178408 Info](#))

Dr. Dietz used to teach at Purdue, and I had the good fortune to take a compiler course taught by him. On the first day, when introducing himself, he came to the part where he was describing how to get into contact with him. When giving out his phone number (at Purdue, on-campus numbers were 5 digits long) he mentioed that his phone number was "GEEKS". He added, "No, I didn't ask for GEEKS, but when I figured it out, I thought it was pretty cool."

Needless to say, it was a pretty cool course.  
[ [Reply to This](#) | [Parent](#) ]

**Wow!** (Score:1, Flamebait)  
by Squeeze Truck ([xmsho@!SHPAM!NEIN!yahoo.com](mailto:xmsho@!SHPAM!NEIN!yahoo.com)) on Thursday December 27, @01:38AM (#2754194)  
([User #2971 Info](#))

Can you imagine a beowulf cluster of these supercomputers??!?!  
...I'll get me hat...  
[ [Reply to This](#) | [Parent](#) ]

**strange...**(Score:1)  
by altan on Thursday December 27, @01:44AM (#2754201)  
([User #519377 Info](#) | <http://slashdot.org/>)

tis weird... just at 3:00 in the morning i was thinking of a 3-4 boxed p2/128mb ram beowulf cluster instead of a pentium4/640mbram...

- Sections**
- [apache](#)  
Dec 15
  - [askslashdot](#)  
Jan 3  
(7 recent)
  - [books](#)  
Dec 21
  - [bsd](#)  
Jan 3  
(2 recent)
  - [developers](#)  
Jan 3  
(4 recent)
  - [features](#)  
Dec 27
  - [interviews](#)  
Dec 21
  - [radio](#)  
Jun 29

[science](#)Jan 3  
(6 recent)[yro](#)Jan 3  
(2 recent)[ [Reply to This](#) | [Parent](#) ]

**A real supercomputer? Not exactly** (Score:5, Insightful)  
 by [fgodfrey](#) ([fgodfrey@bigw.org](mailto:fgodfrey@bigw.org)) on Thursday December 27, @01:49AM  
 (#2754206)  
 (User #116175 Info | <http://forest.bigw.org/>)

The article would have people believe that all a supercomputer is is a collection of not-quite-modern processors, memory, and an interconnect of some sort. This is simply not the case. If it were, why do many (granted a smaller number than before) people still buy real big iron? The answer is that Cringely's (sp?) collection of processors is not a real supercomputer for the kinds of applications that are associated with traditional machines. Traditional vector supercomputers still have processors that are faster than Pentium 4 class systems. Traditional massively parallel supercomputers (which are the most similar to a cluster) have a number of features not found in your average garage built cluster like a truly low-latency interconnect, gang scheduling of entire jobs, single system image for users/administrators/processes.

Clusters are great for embarassingly parallel applications (ie ones that have threads which don't communicate with each other much. This includes things like SETI@home and batch rendering of images. What they don't compare on is applications that communicate a lot like nuclear physics simulations. This is not to say that that will never change in the future, but for the time being it's still true.

Last, and certainly not least, real supercomputers have memory bandwidth that can match the speed of the processor. A Cray or an SGI Origin has an absolutely massive amount of bandwidth from the processor to local memory compared to a PC. That allwos a traditional supercomputer to actually \*achieve\* the fantastic peak performance numbers. On many applications, the working sets are huge and don't fit in cache so you end up relying on memory being fast. On a PC, it's not and I've heard from sources I consider reliable (though I have no actual numbers to back this up so it may be rumor only) that one large cluster site sees around 10% or less of peak on a cluster for a nuclear physics simulation, whereas, on a vector Cray, you can hit ~80% of peak. This means that the cluster has to be 8 times more powerful and when you start multiplying the costs by 8, they start looking like the same price as a real supercomputer.

So my point is that building a real supercomputer does not mean grabbing a bunch of off-the-shelf components, slapping them together with a decent network and running Beowulf (or a similar product).

[ [Reply to This](#) | [Parent](#) ]

**A real supercomputer? Yes, exactly** (Score:4, Interesting)  
by Multics on Thursday December 27, @09:51AM (#2754722)  
([User #45254 Info](#))

Your comments are true for a 486. They are not true for anything much newer. An IBM SP machine, which owns half of the top 10 on the [top500](#) [[top500.org](#)] list, is basically a commodity parts built system.

Yes, these systems are not sometimes the best for handling vectorizable jobs, but they are so inexpensive compared to the old specialized hardware that it is easier to waste cycles than build special hardware.

As to memory bandwidth. Modern CPU caches make the question nearly moot.

If all of this were not true, then people wouldn't be building clusters and the majority of the top500 list wouldn't be dominated by clusters. Instead there are 3 traditional architecture machines in the top 20. This is the reason that Cray (etal) no longer dominates the marketplace... commodity systems have overtaken nearly all of the specialized hardware world.

-- Multics

[ [Reply to This](#) | [Parent](#) ]

**Re:A real supercomputer? Yes, exactly** (Score:2)  
by foobar104 on Thursday December 27, @12:54PM  
(#2755279)  
([User #206452 Info](#))

*As to memory bandwidth. Modern CPU caches make the question nearly moot.*

This is simply not true. Your other points are pretty wacked, too, but I'll take this one because I have personal experience.

I have some image processing code that runs on IRIX, and I recently did a shoot-out between an Origin 2000 and an Origin 3000. Both machines had eight 400 MHz R12000 processors with 8 MB of secondary cache and 4 GB of RAM, and both were equivalently equipped for disk.

The Origin 3000 was almost *twice* as fast as the 2000 was, with identical CPUs, memory, and disk. (The actual numbers are on a spreadsheet at the office, unfortunately.) The

difference? Memory and interprocessor bandwidth. The Origin 3000 platform has a specified memory bandwidth of about 2.5 times the bandwidth of the Origin 2000.

The test involved taking a big multispectral image, splitting it up into tiles, handing each tile off to a thread, and doing some processing on the tiles. The data set was pretty huge, but not so big that it couldn't be cached entirely in RAM, so the first step was to load the whole thing into memory. But for the actual test run, there was a lot of fetch-operate-fetch, which really exercised the memory bandwidth of the system.

So your comment about memory bandwidth being moot is completely off base.

[ [Reply to This](#) | [Parent](#) ]

**Re:Use DDR RAM**(Score:2)

by foobar104 on Thursday December 27, @04:13PM  
(#2756037)

([User #206452 Info](#))

I guess your ideas are kind of on the right track, but you should probably familiarize yourself with modern system architecture trends.

Crossbar-style system interconnects are not new ideas. I'm not an authority on the subject, but I know that the Cray Y-MP had a 32-port switch architecture that provided about 1.3 GB per second of memory bandwidth per processor (hope I'm remembering these numbers right!)

The DEC VAX 9000 series had a 1 GB/second CPU-to-memory pathway that utilized a crossbar switch, also.

Both of these systems were in wide use around 1990, give or take a few years. And, of course, the ideas go back much further than that. I used to have a copy of a paper by Wulf in *Communications of the ACM* dated 1974 that described a switch-based multiprocessor system. Can't find it right now, alas.

Things have come a long way. From 1 GB/sec aggregate in 1990 to 22 GB/sec aggregate in 1998 (the Cray SV1) to 40 GB/sec aggregate in 2001 (the

SV1ex). The SV1ex provides *each processor* with 6.4 GB/sec of bandwidth into and out of main memory.

Increasing the speed of the RAM isn't the issue-- the SV1ex uses commodity SDRAM. The issue is building sufficiently large parallel paths for the memory controllers to execute very large parallel fetches into a vector cache.

So I guess you could say that you're headed in the right direction, but you've got a long way to go. ;-)

[ [Reply to This](#) | [Parent](#) ]

- [1 reply](#) beneath your current threshold.

**Re: A real supercomputer? Yes, exactly!** (Score:2)  
 by [fgodfrey](#) ([fgodfrey@bigw.org](mailto:fgodfrey@bigw.org)) on Thursday December 27,  
 @05:03PM (#2756288)  
 (User #116175 [Info](#) | <http://forest.bigw.org/>)

No longer dominates the Top 500 and no longer dominates the marketplace are two different things. The Top500 benchmark (LINPAC) doesn't do a lot of interprocessor communication and hence is the type of job well suited to a cluster.

As for cycles wasted/cost, that is going to depend on the applications involved. At some point, the sheer cost of the power wasted is going to be a factor. Obviously not on a garage built six node cluster, but if you start talking about 2048p the power *\*will\** be an issue.

The IBM SP, while being *\*mostly\** commodity uses some non-commodity parts and has a lot of proprietary software to make it work.

CPU caches, modern or otherwise, are not an issue with an application that has, say, a 1 gigabyte working set. It simply doesn't fit in the cache no matter what you do. You can restructure loops to make things better, but you're still going to be banging on memory.

You're right, commodity systems have overtaken a lot of areas that used to require traditional supercomputers, but

then, the market for traditional-architecture supercomputers has *\*never\** been big.

[ [Reply to This](#) | [Parent](#) ]

**Re:A real supercomputer? Not exactly** (Score:1)

by [The Fun Guy](#) ([bniemira@arserrc.gov](mailto:bniemira@arserrc.gov)) on Thursday December 27, @ 12:03PM (#2755118)

([User #21791 Info](#) | <http://members.tripod.com/bniemira>)

You said: "The answer is that Cringely's (sp?) collection of processors is not a real supercomputer for the kinds of applications that are associated with traditional machines. ... Clusters are great for embarassingly parallel applications (ie ones that have threads which don't communicate with each other much. This includes things like SETI@home and batch rendering of images."

Cringely said: "Beyond using it to heat my office, I plan to keep the supercomputer busy with a video compression project I'm doing as well as further experiments in wireless communication."

Sounds like he's using the right tool for the job, then. I think he's using the term "supercomputer" to refer to a machine that is many, many times more powerful than the PC that a typical user would have sitting on his/her desk. By this definition, my PIII-600 is a supercomputer compared to the 486SX-25 I started with in 1992.

[ [Reply to This](#) | [Parent](#) ]

**Re:A real supercomputer? - what about software?** (Score:1)

by [udittmer](#) on Thursday December 27, @ 12:35PM (#2755228)

([User #89588 Info](#) | <http://slashdot.org/>)

Lets also not forget the little matter of needing software that can take advantage of parallelism. Not being familiar with QNX I don't know if it can schedule threads or processes across networked computers (my guess would be no), but tools for parallel execution (especially compilers and efficient networking libraries) are still hard to write and generally expensive. Something like MPI can probably be gotten for free these days, but then you're back coding Fortran or C or something, always with a watchful eye on parallel performance. That takes experience and time, both of which aren't cheap. Also, if any communication needs to hapen between processors, network latency usually tends to be a bigger problem than network bandwidth, unless large amounts of data are transferred (in which case good performance may be impossible due to a low CPU/communication ratio).

[ [Reply to This](#) | [Parent](#) ]

**Re:A real supercomputer? Well, depends**(Score:1)  
by Zilya on Thursday December 27, @03:08PM ([#2755769](#))  
([User #228271](#) [Info](#))

Fast ethernet is good enough for pretty big range of well-written parallel applications. In general, the ratio of CPU power/communication has to be proportional to nonlinearity/nonlocality ratio of the problem the computer is designed for. Usually, for partial differential equations, explicit timestepping algorithms put little strain on communications. OTOH, once you go to implicit transport equations of any kind, you need all communications you can get.

And, memory bandwidth is a bitch, I agree. Those dual athlons really have memory bottleneck :(

[ [Reply to This](#) | [Parent](#) ]

**I'll be going with MOSIX**(Score:1)  
by foqn1bo on Thursday December 27, @01:51AM ([#2754212](#))  
([User #519064](#) [Info](#))

A housemate of mine and I decided that we wanted to build a pathetic little Supercomputer out of the various PCs laying around in our little Geek House. We've decided to give MOSIX [[mosix.org](http://mosix.org)] a run. It sounds like a fairy tale solution...especially when it comes to automatic process migration node to node. Anybody here have any positive experiences or harsh words regarding this?

[ [Reply to This](#) | [Parent](#) ]

**:(** (Score:1)  
by NiftyNews on Thursday December 27, @01:53AM ([#2754215](#))  
([User #537829](#) [Info](#) | <http://www.niftynews.com/nndd/main.shtml>)

While he's giving things away...

...can I have a pony?

[ [Reply to This](#) | [Parent](#) ]

**D-Link sells Gigabit NICs for cheap**(Score:1)  
by spullara ([spullara@yahoo.com](mailto:spullara@yahoo.com)) on Thursday December 27, @02:25AM  
([#2754238](#))  
([User #119312](#) [Info](#) | <http://www.sampullara.com/>)

They should recalculate all this if they have Gigabit LAN cards. You can get them from buy.com for about \$57. I have used them at home to make really fast point-to-point links. Also, even though the switches with all gigabit ports start at around \$700 for \$200 you can get a switch with 8 10/100 ports and 1 gigabit port. That should add some interesting properties to the network.

[ [Reply to This](#) | [Parent](#) ]

**Re:D-Link sells Gigabit NICs for cheap** (Score:2)  
by [biglig2](#) on Thursday December 27, @06:22AM (#2754447)  
(User #89374 Info | <http://www.bigwig.net/biglig/> | Last Journal: [Tuesday December 18, @12:23PM](#))

Gb Uplink ports wouldn't really help - the traffic pattern inside KLAT2 is flat, where all nodes are equal. Not like a LAN where it helps to not have a bottleneck at the switch interconnects.

[ [Reply to This](#) | [Parent](#) ]

**Wouldn't help?** (Score:2)  
by [Svartalf](#) ([fearl@!spammers!die!airmail.net](mailto:fearl@!spammers!die!airmail.net)) on Thursday December 27, @11:13AM (#2754945)  
(User #2997 Info | <http://members.nbci.com/svartalf>)

It wouldn't help so long as the CPUs couldn't utilize the gigabit bandwidth. Swap out the 100 megabit lan cards for the cheapo gigabit ones for slightly more money- I think you'll find that this cluster's still starved for bandwidth.

[ [Reply to This](#) | [Parent](#) ]

**Re:D-Link sells Gigabit NICs for cheap** (Score:2)  
by [hey!](#) ([mattleo@treehouse.acrcorp.com](mailto:mattleo@treehouse.acrcorp.com)) on Thursday December 27, @09:03AM (#2754610)  
(User #33014 Info | <http://slashdot.org/>)

How fast are those cards actually?

I know D-Link's PCMCIA 100BaseTX cards are 16-bit, so while they will signal at 100MB/sec, their throughput is not any more than (as far as I can see) than you would get from an old desktop NE2000 adapter. Low end network hardware frequently pulls this kind of stunt -- repackage old technology so that it will look like it should perform better than it actually can.

[ [Reply to This](#) | [Parent](#) ]

- [1 reply](#) beneath your current threshold.
- [1 reply](#) beneath your current threshold.

**"Ultra Wide Band" - no** (Score:2)  
by [Animats](#) ([slashdot-replies@downside.com](mailto:slashdot-replies@downside.com)) on Thursday December 27, @02:29AM (#2754245)  
(User #122034 Info | <http://www.animats.com>)



This is just spread spectrum, but with even more spread. See [TimeDomain \[timedomain.com\]](http://timedomain.com) for the hype. Even they admit "*UWB's best applications are for indoor use in high-clutter environments.* We already have wireless LANs, and they work quite well. UWB may or may not play in that market, but it's not a big deal.

The FCC is being very cautious about mass-market UWB products. Since these things blither over a gigahertz or so of spectrum, they overlap with other services. At low power, a few of these things are probably OK, but in bulk, there could be trouble. The concern is that mass deployment could wipe out other services in congested areas.

[ [Reply to This](#) | [Parent](#) ]

**Re:"Ultra Wide Band" - noise**(Score:1)

by Student\_Tech on Thursday December 27, @03:11AM

(#2754282)

([User #66719 Info](#))

From my understanding, and my dad's, they also raise the noise floor so if you are trying to talk to people over radio (such as amateur radio) it is that much more noise you must contend with. Yes, DSPs can help to some extent, but UWB still raises the amount of noise one must deal with.

[ [Reply to This](#) | [Parent](#) ]

**Re:"Ultra Wide Band" - noise**(Score:1)

by markov\_chain on Thursday December 27, @03:45PM

(#2755926)

([User #202465 Info](#))

Right, that's essentially what the previous post said. The UWB transmissions occur in a very wide band of frequencies, thus overlapping with other services like amateur radio. To the other services any power not coming from the original transmitter is counted as noise; that's why UWB transmissions are said to raise the noise floor.

The argument for UWB is that a single UWB transmitter's power is spread so much that it adds only tiny amounts of noise to other services. However, the question is what happens when there are massive amounts of UWB transmitters contributing noise.

[ [Reply to This](#) | [Parent](#) ]

**Sorry had to point this out**(Score:1, Informative)

by Anonymous Coward on Thursday December 27, @02:49AM

(#2754267)

From the PBS article:

"The solution was to put more cheap Ethernet cards in each PC, and then use "channel bonding" to make them all look like a single faster card"

From KLAT2 FAQ

"Every NIC in every PC has a unique MAC address (and potentially unique IP address) -- i.e., this is not channel bonding."

FNN is totally different and in many cases more suited for this app than simple channel bonding.

One thing I did wonder about was. Why the floppy drive? You can get netboot cards very cheap... And you'd only need one per system. Just one less mechanical thing to fail. Plus the node would come up much faster. PXE or even BOOTP/DHCP boot would be fine.

Also I kind of wonder about commodity Realtek cards. I'm sure Realtek makes a fine chipset, but most vendors who use Realtek chipsets really skimp on the rest of the card. You can get 3com or Intel Pro/100 multipacks almost for the same price as the realtek cards sell for off of the shelf.

Ok rant over... Flame away.

[ [Reply to This](#) | [Parent](#) ]

**Not really here**(Score:1)

by [WyldOne](#) ([{cwyles}](#) [{at}](#) [{nyx.net}](#)) on Thursday December 27, @02:26PM (#2755567)

([User #29955 Info](#) | <http://www.geocities.com/cwyles/>)

I can get a clone Realtek card for \$10 (or less sometimes) retail. Most other cards are at least \$20. I can also pick up a floppy for \$1 a piece used.

Not sure how much you could get the realteks w/rom for bulk but you would have to call around

[ [Reply to This](#) | [Parent](#) ]

**Re:Not really here**(Score:1)

by [sh4na](#) ([.shana. .at. .esoterica.pt.](#)) on Friday December 28, @05:45AM (#2758033)

([User #107124 Info](#) | <http://w3.to/shana>)

Sure you can get any crappy card for 2 cents, but then you'll get the quality of a 2 cent card. I've tried Realteks and the likes, and in a 100Mbit network I'd get the performance of a 10, sometimes worse. It took ages to do anything! So no thank you, I'd rather stick with my \$20 card.

Floppy drives used? Unless it's more than 5 years old (when they used to make them good), it'll last you through two or three bad diskettes, it'll read half of the good ones, and it'll smash the rest. Another good use for your money, I'd say...

By the way, I can sell you a really good used HD, cheap. If you don't mind a few bad blocks here and there, that is...

You know, sometimes you get what you pay for.

[ [Reply to This](#) | [Parent](#) ]

**links to really homemade supercomputers** (Score: 1)  
by DrD8m on Thursday December 27, @03:31AM ([#2754299](#))  
(User [#307736](#) Info | <http://www.sorgonet.com>)

I like this kind of home supercomputers [some examples here](#) [sorgonet.com] of course, linux powered

[ [Reply to This](#) | [Parent](#) ]

**Better throughput than Gigabit** (Score: 1)  
by NanoGator on Thursday December 27, @04:13AM ([#2754335](#))  
(User [#522640](#) Info | <http://www.nanogator.com>)

How can 4 bonded 10/100 cards provide better throughput than a gigabit card? I can understand them being cheaper, but the idea of them being faster is hard for me to grasp.

The way I see it, 4 100 megabit cards at MOST would create a 400 megabit pipe. Wouldn't bonding add overhead to it to make it even less than 400 megabits?

Would somebody mind explaining how this works?

[ [Reply to This](#) | [Parent](#) ]

**Re:Better throughput than Gigabit** (Score:1)

by Defiler on Thursday December 27, @10:01AM (#2754746)

[\(User #1693 Info\)](#)

Simple. No one actually gets 1000mbit from a GigE card. You're lucky to get 50MB/sec with a crossover cable. You can match that with four 100Base-TX cards. Actually, GigE cards are extremely cheap. Only about double the price of cheap 100mbit cards. It's the switches that are expensive. Still at least \$100 per port.

[ [Reply to This](#) | [Parent](#) ]**Re:Better throughput than Gigabit** (Score:1)

by Extreme Unguent on Thursday December 27, @12:41PM

(#2755250)

[\(User #546398 Info\)](#)

Well, KLAT2 isn't using channel-bonding and they are talking about bisectional b/w not local... so it's not a very adaptable comparison. He also has custom (Realtek) drivers for his NICs.

[ [Reply to This](#) | [Parent](#) ]**Re:Better throughput than Gigabit** (Score:1)

by Defiler on Thursday December 27, @10:49PM

(#2757288)

[\(User #1693 Info\)](#)

Sure. I didn't say that was what KLAT2 was using, just making the point that you don't actually see a full gigabit with GigE equipment.

[ [Reply to This](#) | [Parent](#) ]**Re:Better throughput than Gigabit** (Score:1)by [WyldOne](#) ([@cwyles](#)) ([@nyx.net](#)) on Thursday December 27,

@02:10PM (#2755504)

[\(User #29955 Info\)](#) | <http://www.geocities.com/cwyles/>

Ethernet works by the dump the data on the wire and hope for no collisions.

Every time you get a collision 2 pcs have to resend their data. so you reduce the bandwidth by 66%.

So by adding more channels to send on you decrease collisions and increase throughput.

[ [Reply to This](#) | [Parent](#) ]

- [1 reply](#) beneath your current threshold.

**Disappointing**(Score:1)

by [MattGWU](#) ([MattSlash at bogonflux dot net](#)) on Thursday December 27, @05:01AM (#2754373)  
 (User #86623 [Info](#))

When I read the title, I had visions of actual homebrew \*supercomputers\*...something along the lines of Euclid from Pi. Break out the soldering irons and damn the torpedoes! Yes, Beowulf technology is great. Genetic algorithms and channel bonding and QNX are nice touches, but Beowulf clusters are fairly common, even for ordinary people...an article about NEW uses for the things would have been nice, other than that, just shut up and build one! Soon to be taking my own advice...even have an app in the works for it. Still a nice project with a half-decent writeup, but it's been done.

How about it, folks? Homebrew big iron? Where would one even begin? Food for thought, at least.

[ [Reply to This](#) | [Parent](#) ]

**Supercomputer in every neighborhood**(Score:1)

by [tempmpi](#) on Thursday December 27, @06:08AM (#2754437)  
 (User #233132 [Info](#))

A supercomputer in every garage is something that is too expensive and useless most of the time, but a supercomputer in the neighborhood could be a realistic and useful idea. We have seen on slashdot that it is possible and not too expensive to make a neighborhood fiber lan. Most people do not need much processing power all the time they need a high peak processing power for short moments.

Most user application software will not be able to use divided the load over the cluster but there would be many applications running at a time, so the load would be spread over the cluster without special application software. People could keep their old PCs and turn them into X terminals or use vnc to connect to the cluster.

\$41,000 is only \$205 for everyone if 200 people use the cluster.

[ [Reply to This](#) | [Parent](#) ]

**QNX? Hey Cringely..**(Score:3, Funny)

by [Chazmati](#) on Thursday December 27, @07:25AM (#2754479)  
 (User #214538 [Info](#))

*"(the operating system) will be QNX, a real time OS that supports massive parallelism and has very low overhead. QNX is fast! QNX is also Posix compliant, so there is lots of software that almost works under it."*

If you're looking for software that **almost** works, I know of an OS that might fit your needs. You're not going to hook this thing up to the Internet, though, are you?

[ [Reply to This](#) | [Parent](#) ]

**Re:QNX? Hey Cringely..(Score:1)**

by Unknown Bovine Group on Thursday December 27, @09:47AM (#2754713)

(User #462144 Info | <http://slashdot.org/>)

Even more interesting... *And even though QNX is a commercial operating system, it is free for noncommercial purposes like mine.*

Uh, Cringely, wouldn't creating the thing and then using it as the subject of an article for the company that employs you count as a commercial purpose?

[ [Reply to This](#) | [Parent](#) ]

**Re:QNX? Hey Cringely..(Score:2)**

by Knobby on Thursday December 27, @04:13PM

(#2756040)

(User #71829 Info)

*Uh, Cringely, wouldn't creating the thing and then using it as the subject of an article for the company that employs you count as a commercial purpose?*

You don't really expect QNX to bitch about a little free advertising do you?..

[ [Reply to This](#) | [Parent](#) ]

**only fastest one percent is supercomputers(Score:3, Informative)**

by peter303 on Thursday December 27, @09:43AM (#2754698)

(User #12292 Info)

By definition only the fastest devices are supercomputers. These days that is about a teraflop. Ththa includes the US DOE ASCI series and the announced installation of the Blue Storm and Blue Gene IBM computers. Ten gigaflop computers a dime a dozen and a hundred gigaflops not so rare.

[ [Reply to This](#) | [Parent](#) ]

**Cringely Wants A Supercomputer in Every Garage(Score:1)**

by ZaneMcAuley on Thursday December 27, @11:36AM (#2755029)

(User #266747 Info | <http://www.ineedanewgirlfriend.com/> | Last Journal: [Thursday October 25, @04:32PM](#))

Well, don't we all :)

[ [Reply to This](#) | [Parent](#) ]

**This is all OBE (Overcome by events)(Score:1)**

by Extreme Unguent on Thursday December 27, @11:51AM (#2755077)

(User #546398 Info)

Most of what Cringely said may have been true a year or so ago but it isn't now. For instance, if he is using Athlon XPs, why would he move FP code to 3DNow! instead of SSE? And there are a host of competing cheap interconnects now, especially if you can avoid TCP/IP. But if you can't, there's always IP-over-FireWire... M\$ has had that running at 400Mbps for years. I don't think you can get into UWB for \$6000. See [the Linux Clustering Info Center \[lcic.org\]](#) and [Extreme-Linux.com \[opnsrc.com\]](#).

[ [Reply to This](#) | [Parent](#) ]

**Why spend \$40,000 when you only need \$13,000?** (Score: 1, Informative)  
by Anonymous Coward on Thursday December 27, @01:30PM  
([#2755375](#))

I don't understand why this guy's so excited about 64 gigaflops for \$40,000. Consider the following

6 Dual 800Mhz G4 PowerMacs (running your choice of Darwin or Linux)  
Processing Power: 70.8 gigaflops (11.8 each machine)  
RAM included: 1.5 GB (heck, I already have 1GB in my PowerBook G4)

Hard Disk Space: 480 GB

Extras: Each machine has one 4x AGP slot, four 64-bit PCI slots, an NVIDIA GeForce 2 MX card (64MB), Gigabit Ethernet and a 56K modem you can just take out and hang on your Christmas Tree or something

Cost: \$21,000 (\$3,500 from store.apple.com)

Ebay off the CD-R/DVD-R Drivers for \$350 each to save \$2,100

Total cost: \$18,900

Tip: If you wait until after MacWorld in January (where Apple will most likely introduce a line of faster PowerMacs) you can either get more gigaflops for the same price or lower the price of 64 gigaflops to around \$13,000.

What kind of an idiot goes out and spends over \$40,000 on something less powerful than he could get for \$13,000? And people complain that Apple hardware is expensive... morons...

[ [Reply to This](#) | [Parent](#) ]

**Re: Why spend \$40,000 when you only need \$13,000** (Score:1)  
 by Hankd on Saturday December 29, @09:30AM (#2762106)  
 (User #166718 Info)

KLAT2 is nearly 2 years old and it gets 65GFLOPS on a real application... peak speed is 180GFLOPS. The G4 doesn't do badly speed-wise, but is not price/performance competitive with Athlons... in fact, right no general-purpose processor is. ;-)

[ [Reply to This](#) | [Parent](#) ]

- [1 reply](#) beneath your current threshold.

**I want a supercomputer** (Score:1)  
 by RMBWebmaster on Thursday December 27, @02:36PM (#2755612)  
 (User #542302 Info | <http://www.ratemybuds.com/>)

Can this be done with sparc's?

[ [Reply to This](#) | [Parent](#) ]

**Uses for mini-cluster** (Score:2, Funny)  
 by WyldOne ( {cwyles} {at} {nyx.net} ) on Thursday December 27, @02:41PM  
 (#2755642)  
 (User #29955 Info | <http://www.geocities.com/cwyles/>)

- 1) heat in garage in winter
- 2) Top 10 in Seti@home
- 3) Porno-ize you favorite anime (Final Fantasy anyone?)
- 4) Why are you reading this? I thought you were doing #3

[ [Reply to This](#) | [Parent](#) ]

**Thanks for reading the articles before posting** (Score:1)  
 by VonSnaggle on Thursday December 27, @03:59PM (#2755978)  
 (User #64586 Info)

Not to sound like an ass, but this is the best article I've read on Slashdot for quite some time. So when I want more information on the subject I was hoping to find some intelligent or at least funny conversation and most of the posts are crap, just when my 3 day moderating period has ended. I wish some people would just read the linked articles before posting!

[ [Reply to This](#) | [Parent](#) ]

**Essential Costs** (Score:1)  
 by Tekgno on Friday December 28, @12:43AM (#2757614)  
 (User #321071 Info)

When you are comparing costs [[aggregate.org](http://aggregate.org)] for construction, don't forget the essentials.

16 Pizzas for student helpers @ \$10  
 4 Cases of soda student helpers @ \$7

Total: \$188



[ [Reply to This](#) | [Parent](#) ]

### **Supercomputers and Bisection Bandwidth** (Score:1)

by Hankd on Friday December 28, @09:56AM ([#2758475](#))

([User #166718](#) [Info](#))

Two useful definitions that explain why KLAT2 was built as it was:

- A *supercomputer* is a computer that is not only very fast, but whose design allows it to be scaled-up to be faster as more money is spent.
- *Bisection Bandwidth*, the worst-case total bandwidth between halves of a parallel machine when all processors are communicating, is the primary measure of supercomputer network bandwidth, **NOT** NIC speed. Further, NIC performance is often limited by the OS interface and/or PCI bus. This is why a network made of multiple 100Mb/s NICs per PC and cheap wire-speed switches easily can equal or exceed the performance of using Gb/s NICs and the narrower, often less than wire speed, Gb/s switches.  
The same argument applies for latency: single switch for 100Mb/s FNN versus multiple switch hops for Gb/s.

[ [Reply to This](#) | [Parent](#) ]

### **There is no one formula to build a cluster** (Score:1)

by [mikefoley](#) ([mike.foley@telogist.com](mailto:mike.foley@telogist.com) [['chno' in gap](#)]) on Friday December 28, @10:25PM ([#2761386](#))

([User #51521](#) [Info](#) | <http://www.yelof.com/>)

It's interesting reading the replies here. It's as if all clusters are only defined by CPU and speed of the NIC. Sorry, it's not that way at all.

A cluster should be designed to solve a specific problem. You have to do some math up front before making your choices on things like NIC's and CPU's. You can't apply a blanket solution. Some of the things to consider are:

- . The size of the data being exchange
- . Does the matrix you are computing fit into cache?
- . Is the problem dependant on memory bandwidth?
- . etc....

For example, for some problems, the amount of data exchanged fits well within a Jumbo Packet of a Gb Ethernet AND is not affected by latency.

Other problems are very dependant on latency and require things like a Dolphin or Myrinet card. These are not inexpensive items.

Still other problems require memory bandwidth and work well with systems like Alpha's.

So, if you are building a small cluster to run POVray at home, go with cheap. If you are trying to crack the human genome, you need racks upon racks of things like high-end Alpha's with Quadrics interconnects. It's all dependant on the problem you are trying to solve.

FWIW, I used to work at API and DEC/Compaq.

[ [Reply to This](#) | [Parent](#) ]

**There are design rules and cluster design tools** (Score:1)  
by Hankd on Saturday December 29, @09:24AM (#2762102)  
(User #166718 [Info](#))

There is no formula, but there are lots of design rules... check out:

<http://aggregate.org/CDR/the Cluster Design Rules tool>  
[[aggregate.org](http://aggregate.org)]

You specify some characteristics of your application, your site (power and space), and budget; it presents the best designs taken from a design space of millions.

[ [Reply to This](#) | [Parent](#) ]

**a few thoughts** (Score:1)  
by [spiffy\\_guy](#) ([jhs97j@timon.acuFORTRAN.edu](mailto:jhs97j@timon.acuFORTRAN.edu) minus language) on Saturday December 29, @12:33AM (#2761544)  
(User #30225 [Info](#) | <http://slashdot.org/>)

I think one of the neatest ideas about having a supercomputer in your house is that there is nothing to do with it. Most people think this is a negative, having no problems to solve. I think it's fantabulous. Think of all the new problems to discover so you can have something to solve. Sure the lazy among us will go for Optimal Golum Rulers, digits of Pi, Chess, SETI, RC5-xxx, video rendering/compression/effects, or whatever.

IBM was going to/is build a monster of a machine Blue Gene. Biggest machine ever built. However they wanted a problem to show off their beast, they decided to look at protien folding. NOBODY was looking at protien folding using supercomputers to actually analyze them. It was a computer waiting for a problem big enough, and it fit.

I'd also like to note that though these computers being built have HUGE processing power they don't have the latency and bandwidth a lot of problems like Weather prediction. Unless you get a Cray, Superdome, or Regatta it's like towing a trailer with a piece of yarn. For certain problems the processing power is less important than the memory bandwidth.

[ [Reply to This](#) | [Parent](#) ]

**Re:And I want \$10M in every pocketbook**(Score:2, Interesting)  
 by innocent\_white\_lamb on Thursday December 27, @ 12:42AM  
 (#2754120)  
 (User #151825 Info)

Back about... 15 years ago? 20? Something like that.

I was in the office of a research company and the owner showed me their shiny new minicomputer. I can't remember what kind it was, unfortunately.

He said something then that struck me as very insightful and I've not forgotten it to this day.

"You know, minicomputers are looking more like micros every day."

[ [Reply to This](#) | [Parent](#) ]

**Re:Stephen King, author, dead at 55**(Score:1)  
 by SpacePunk (sensei@techdojo.net) on Thursday December 27, @03:41PM  
 (#2755915)  
 (User #17960 Info | <http://slashdot.org/>)

I think that so far in the year 2001 Stephen King has died at least four times. It's one of those things that keeps popping up like those damn gophers in the arcade machine where you hit em with a mallet (whatever the hell it's called).

-

[ [Reply to This](#) | [Parent](#) ]

- **1 reply beneath your current threshold.**
- **10 replies beneath your current threshold.**

[1](#) | (2) (Slashdot Overload: CommentLimit 50)



*Ocean: A body of water occupying about two-thirds of a world  
 made for man -- who has no gills. -- Ambrose Bierce*

All trademarks and copyrights on this page are owned by their respective owners. Comments are owned by the Poster.

The Rest © 1997-2001 [OSDN](#).

[ [home](#) | [awards](#) | [contribute story](#) | [older articles](#) | [OSDN](#) | [advertise](#) | [about](#) | [terms of service](#) | [privacy](#) | [faq](#) ]