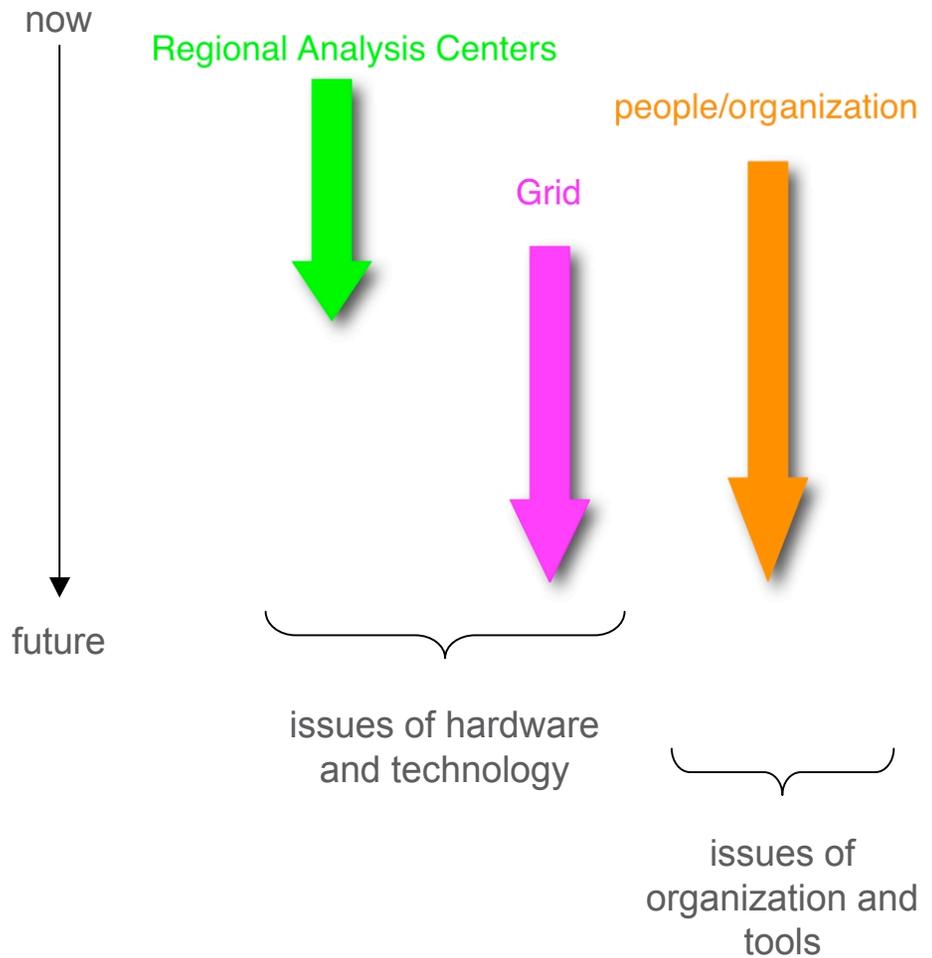


# status, procedure and future strategy for analysis

How do we organize ourselves for worldwide analysis

Like Caesar's Gaul, this talk is divided into three parts:



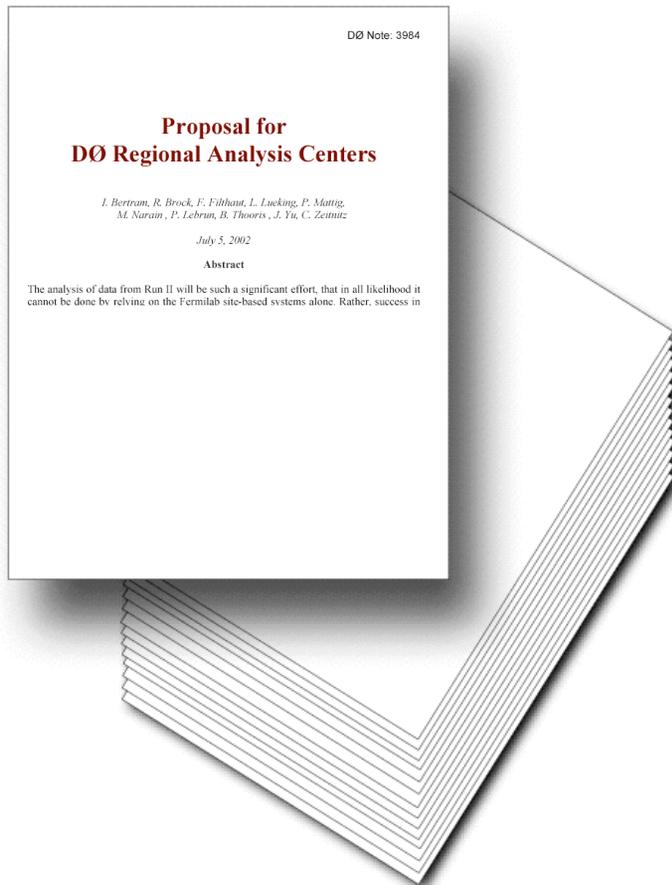
key:

 the purple arrow is to point out items which might benefit from discussion here

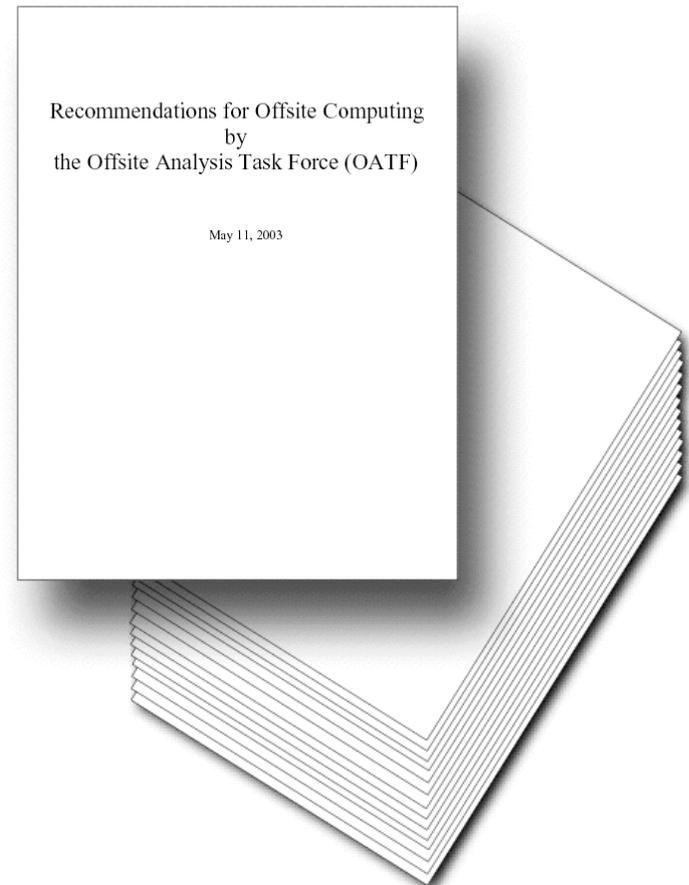
# Regional Analysis Centers



# a system proposed last summer *reviewed last fall/winter (“OATF”)*



Brock, *in absentia*, Beaune Workshop



draft 4: 6/15/03

16 June 2003

**the details:** see *Nick & Christian's talks*

## Tiered, hubbed Computing complexes

### 1. Centralized RAC:

- Caching service: *1/n storage of DST's & at least 1 complete TMB set*
- Database access: database server proxy
- CPU availability: enough to skim/mine data
- Possibly:
  - Reprocessing capability*
  - MC farm capability*

### 2. Related, regional Institutions:

- Institutional Analysis Centers (IAC) - considerable CPU in IAC's
- Possibly
  - Reprocessing capability*
  - MC farm capability*

## Regionalized for

*early deployment efficiency*

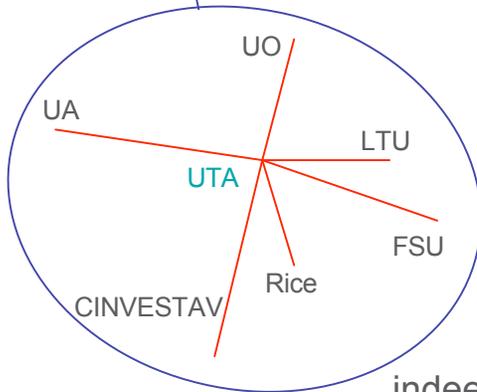
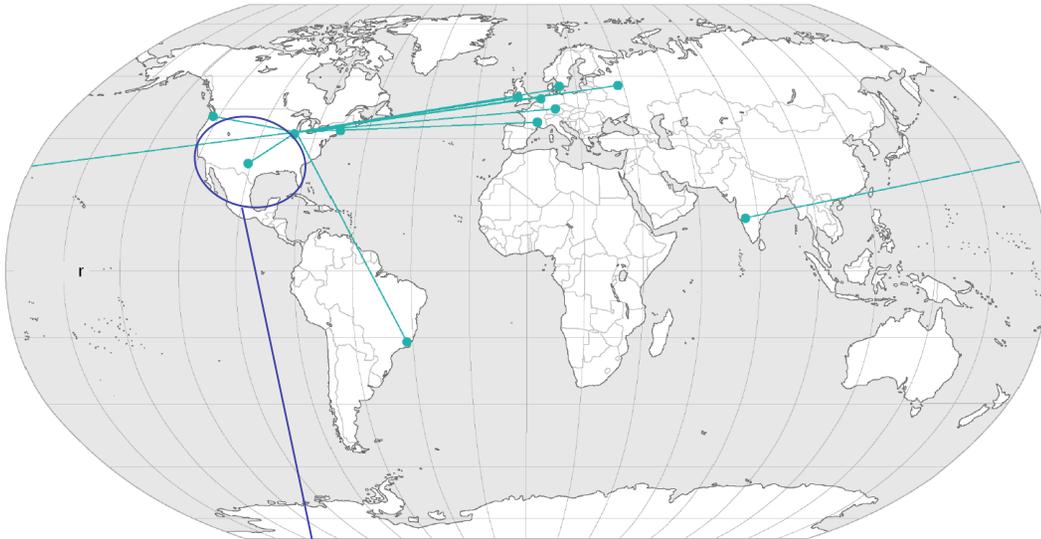
*financial/political leverage*

*in order to train/maintain local support*

- From both IT and Physicist

# the basic scheme:

This (presuming ~10 RAC's):

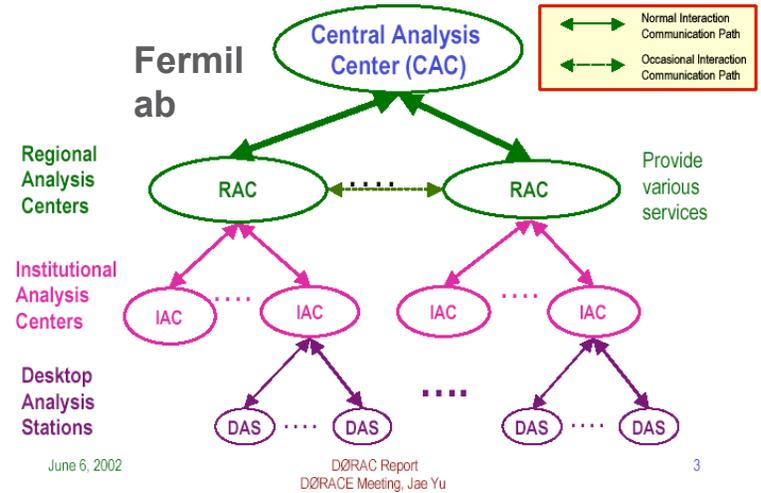


(arbitrary imagined distribution)

already, sponsored a Regional workshop on analysis

indeed, UTA has resources:

## Proposed DØRAM Architecture



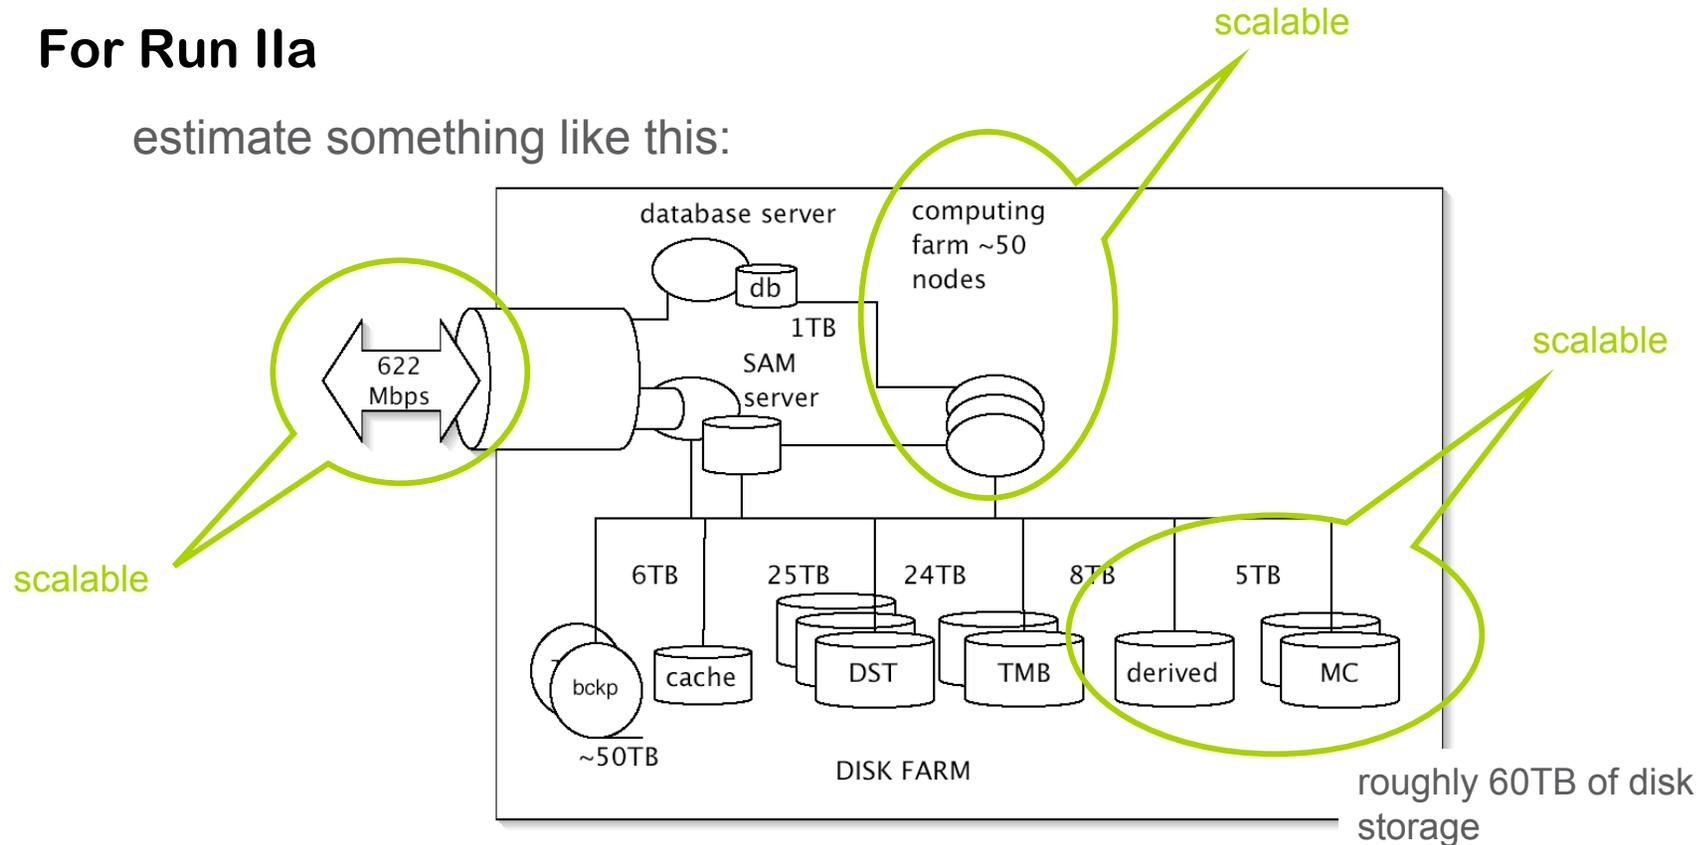
Rather than this:



# summary of an entry-level RAC

## For Run IIa

estimate something like this:



- For 10 RAC's this alone adds > 500 cpu's, deployed in an efficient way - where the physicists are
- IAC's should have have additional, considerable capability
- All in host countries.

# scale it up...to a very significant RAC

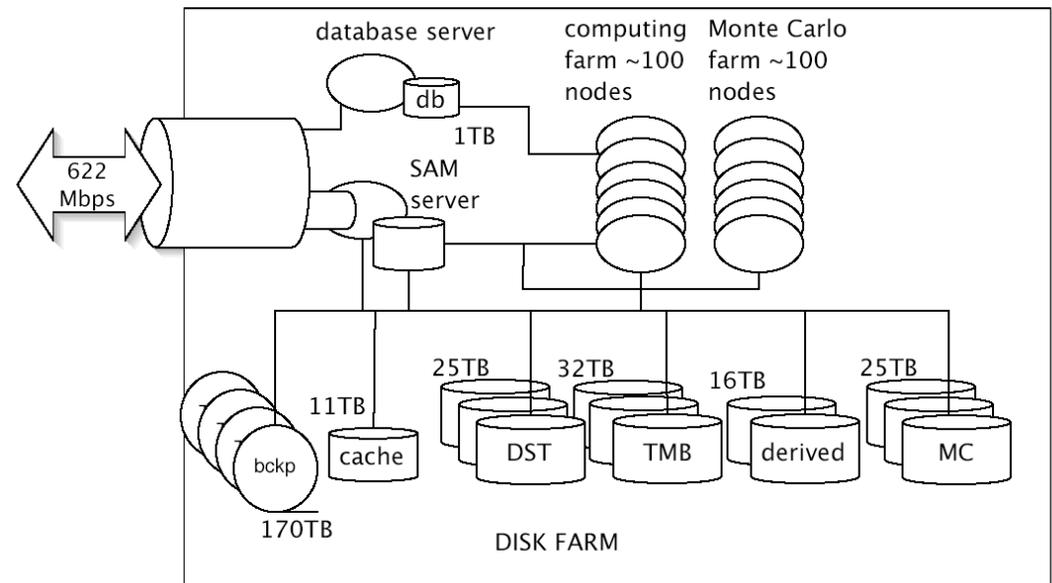
**Keep:** *DST storage as a common resource, SAM, db proxy server*

**Add:**

- More batch computing: +50 nodes
- MC generation: + ~ 100 nodes
- More MC storage: +20TB
- All TMB: +8TB
- More derived data cache: +8TB
- More temporary cache: +5TB

**So:**

*~100TB of disk and ~200 cpu's*



**a very serious system**

desirable to have a few  
may fit as parts of larger facilities

- No longer manageable by a single university department

# Evolution, not revolution

## Near term:

*RAC's have an organizational & development role*

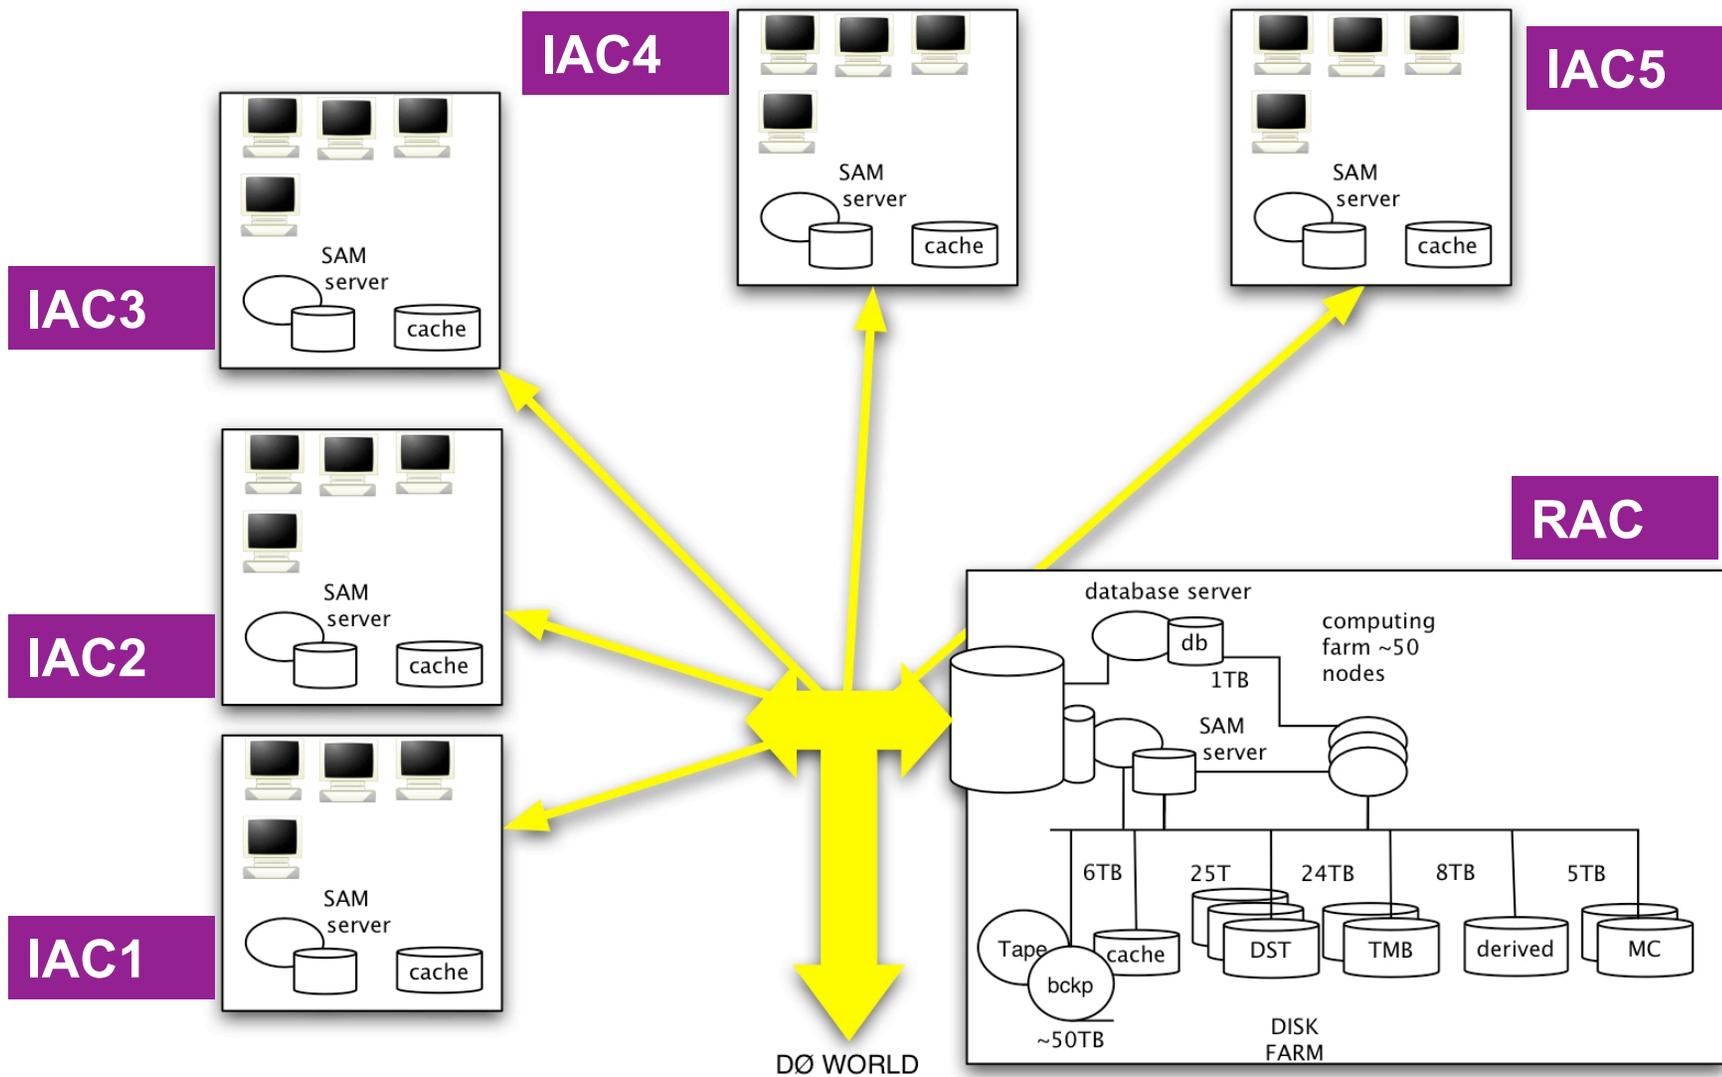
- they relate especially to their Regions as support and source of data
- IAC's mostly “look out” to DØ, in part through their RAC's

## Far term:

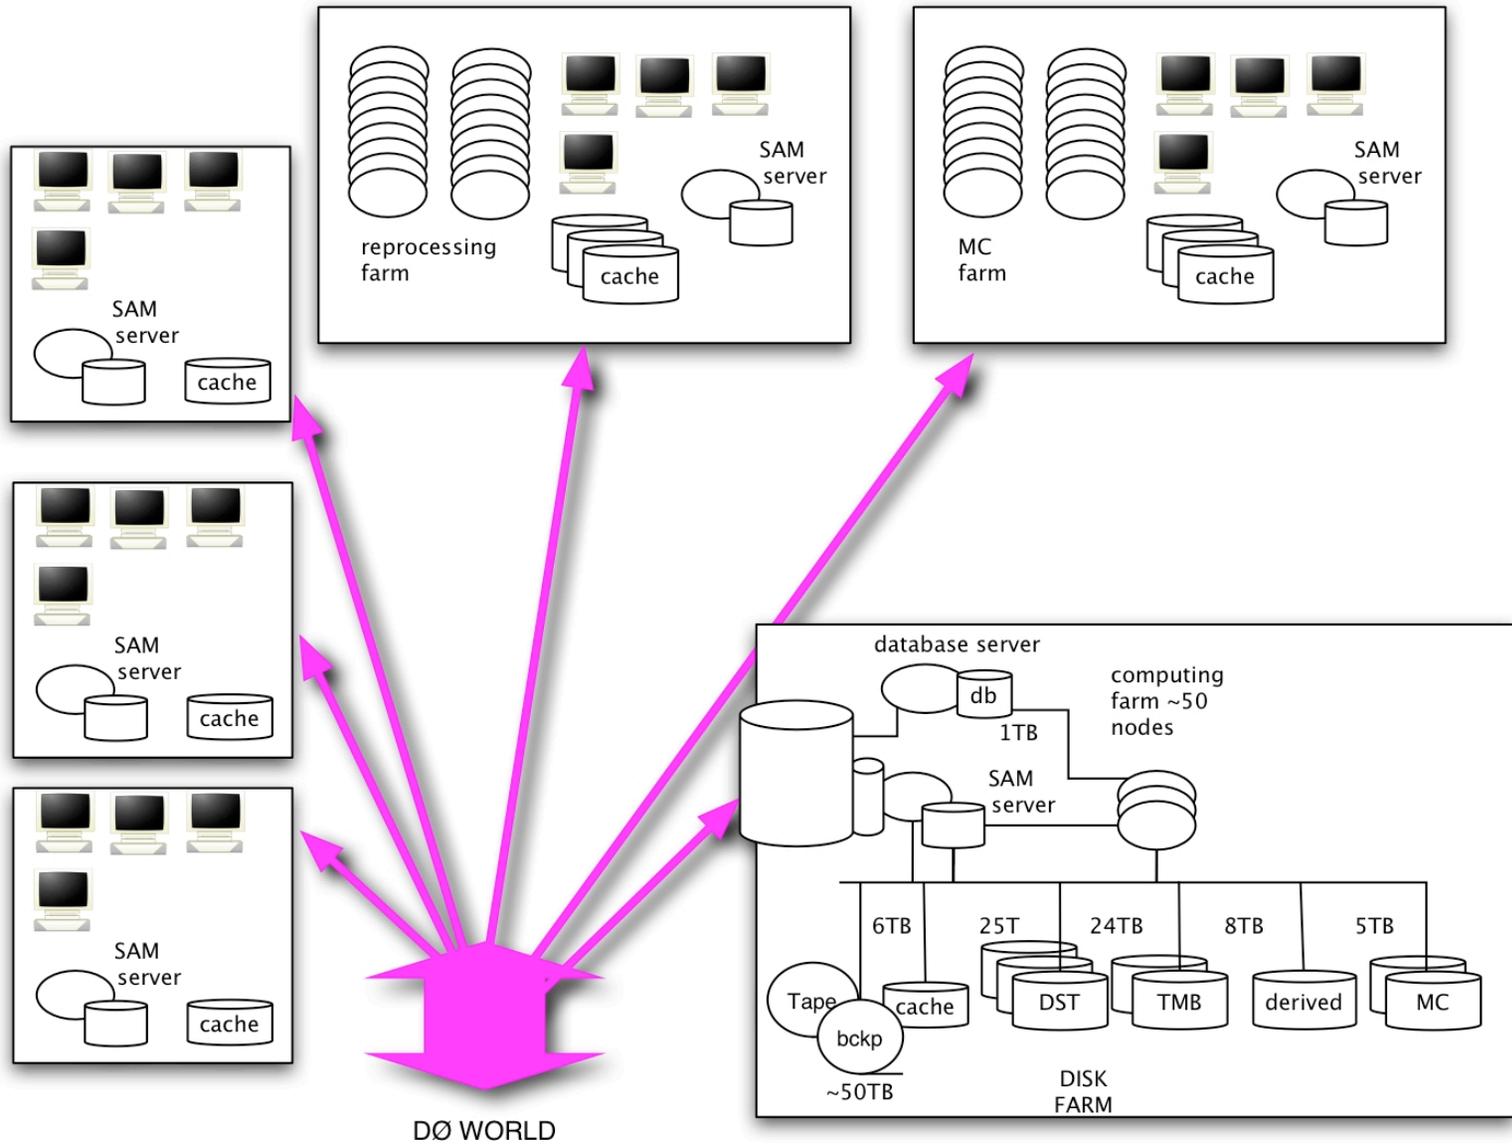
*RAC's and IAC's should evolve:*

- to become components of the whole DØ computing universe  
*all to be come “remote Analysis Centers”*
- while still centralizing expensive caching services  
*with now DØ also “looking in” to RAC and IAC's*
- possibly still serving as point of contact to DØ Support structure

# Go from, initially, a set of capable Regions



# ...to a worldwide system

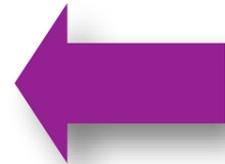


## RAC challenges

- software release management
- specialized FNAL support
- database propagation
- maintain near-real time data (DST/TMB/derived) distribution from FNAL

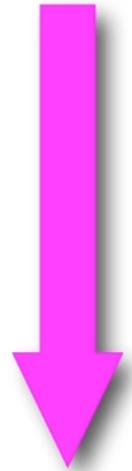
### Beaune:

1. How do we initiate and support RegionalAC's and associated IAC's
2. How do we eventually stage RegionalAC's and IAC's into *remoteAC's*
3. What are the lessons of the prototype Karlsruhe experience?
4. Where will the next RAC/IAC's be located?



This evolution suggests a Grid solution

Grid



## **SAM-Grid/JIM:** see *Wyatt and Kors' talks*

The DØ Grid will be a set of protocols and tools to:

1. *characterize a JOB*
2. *process requests for processing from anywhere in the DØ Virtual Organization (VO)*
3. *seek out, match, and broker for appropriate, available resources within the DØ VO (CAB, RAC's, IAC's, NPACI, etc.)*
4. *scatter the JOB to the collection of willing resource locations*
  - authenticate the user
  - locate, and acquire or remotely use, necessary data
  - securely spawn appropriate processes with remote schedulers
5. *deliver the results back to the requestor and/or catalog any derived data to SAM*

**DØ's version called SAM-Grid or JIM** (Job Information Management)

*It's an ambitious plan*

- Built on SAM
- And it's an advanced R&D effort within the FNAL CD

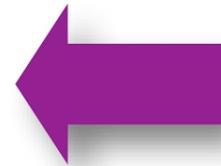
# DØ SAM-Grid Challenges

- are the planned tools the right tools?  
*what do the Physics/ID Groups need? What are they doing now?*
  - new project to document strategies, data formats, tools, etc for each
- how fancy should SAM-Grid seek to be?  
*the “standards” issue: matching SAM-Grid to LCG? (LHC Computing Grid)*
  - FNAL developers are in short supply and great demand
- how do we stage RAC/IAC to SAM-Grid?  
*all at once? inconceivable*  
*horizontally, by institution? maybe...pick a nimble RAC or IAC as □ site*  
*vertically, tool, by tool? e.g., MC brokering/submission first? etc.*
  - we cannot screw up ongoing analysis by being overly aggressive here.
- how do we develop both SAM-Grid for DØ and CDF simultaneously?  
*different analysis models make this complicated*  
*inclusion of CDF into the SAM universe has cost DØ*
- the “D” word: JIM documentation is pretty sparse.

# DØ SAM-Grid Challenges, cont.

## Beaune:

1. How can we stage SAM-Grid deployment into the RAC system?  
Vertically? Horizontally? both?
2. Assess a realistic state of development of SAM-Grid.  
*unrealistic talk of June rollout*
3. Long overdue: establish organization and **oversight** of SAM-Grid project as it pertains to DØ, PPDG, FNAL, etc.  
*of utmost importance: a DØ organizational structure that crosses experiment boundaries to CD, PPDG, LCG, and is woven into the DØ Physics groups*
4. What is a realistic assessment of the development effort to mate SAM-Grid to LCG?

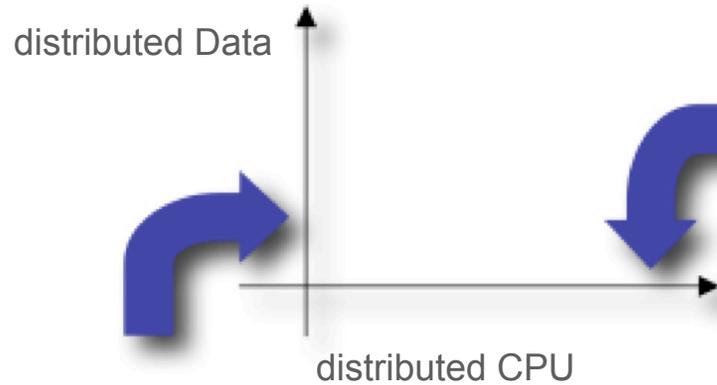


people/organization



# A practical Grid is 3 dimensional

Not just:

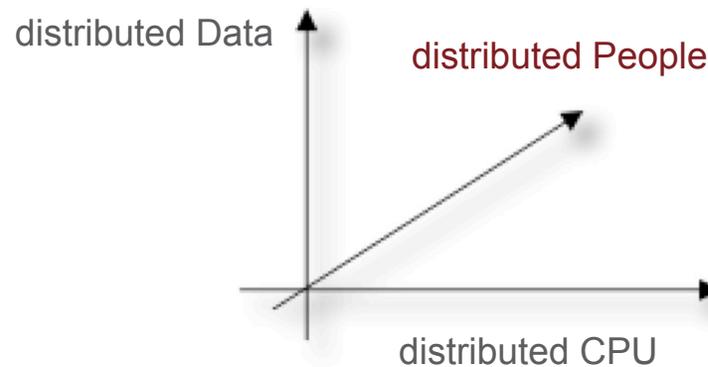


a lot of effort goes into learning to manage resources along this axis - the traditional concentration

SAM, provides us flexibility along this axis, and essentially does it now

But:

I think that this is new, at the scale of even DØ, let alone LHC



# So, does JIM address the right challenges?

Grid discussions are usually hardware. But:



infinite CPU/storage resources do not solve our problems!

## ***The Hard Problems Argument:***

Success in Run II requires heroic analysis efforts:

*multiple fb<sup>-1</sup> of data imply unprecedented burdens:*

- frightening systematic control for coming precision measurements*
- challenging detector calibrations and background control*
- efficient management of enormous primary & derived datasets plus organizing pb's of MC production and reprocessing needs*

Constant of our motion = 200: the number of seats at FNAL

- But, 200 brains is insufficient to solve our intellectual problems*
- we need the full intellectual participation of the entire experiment*
  - 300-400 brains

# I know: tiering is anathema to Grid-aficionados...

But, just as “the best politics is local,”

- *meaning good political organization is personal and interactive*

Physics analysis is also such a human process

- **Old way:** *most work @FNAL, remote people periodically traveled*
- **New way:** *work is distributed @everywhere - as intense off-site as on..*

But, isolated, few-person groups are not likely to be effective

So, an additional focus for RAC's

- **could also be on physics analysis organization:**

for example:

- *Now, senior non-Illinois colleagues without analysis centers have a choice: go nowhere, or travel as much as thousands of miles*
- *But, if RAC concept works, they could have opportunity to be productive traveling a few hundred miles to be with like-minded people*
- *engaging the students before they leave for Illinois and after they return*

RAC's can be a point of contact and support

RAC's could help to empower the entire collaboration

*If you buy the “**Hard Problems Argument**”...the **New Way** is necessary.*

**So, what we need is not  
*Offsite Computing Organization*  
but  
*Offsite Analysis Organization***

# what sort of collaborative work do we do?

## Personal communication

### *Most common current practice:*

- tons of serial email - asynchronous; both parties choose to be involved on their time scale
- office drop-in/hallway encounters - real time; maybe a social penalty...guest interrupts host

## Collaborative Document preparation

### *Most common current practice:*

- MS Word (ugh) or LaTeX source files distributed serially as email attachments. Comments: serial email.

## Desktop sharing

### *Most common current practice:*

- probably VNC, which is a little shaky sometimes, sometimes difficult to all command mouse equally, slow, insecure

## Meetings...we meet a lot.

### *Most common current practice:*

- Conference telephone calls - decent mix with VNC
- Video conferencing, usually ISDN, H.323 IP-based, VRVS  
*video conferencing is still unsatisfactory*

## Virtual shifts

### *Most common current practice:*

- GM remote shifts - web, X based
- Remote SAM shifts - web based

## More?

## Can we do better than some of this?

# we maybe don't need to reinvent this wheel

There is a whole discipline centered around:

- *how distributed teams of people work and how best to create tools to support virtual team collaboration*
- *“Collaboratories” is the buzz word.*
  - LHC has discovered it at some level: Recent ATLAS/CMS ITR; “Opportunities for Use and Development of Collaborative Tools in ATLAS,” atnot /Note /gen-2003-002; much more.
- *LBNL Distributed Collaboratories Project ; UMich School of Information; and many more*

It is largely tools-based

- *what kind of groupware facilitate virtual collaborative efforts?*

We know of some:

- *Video conferencing: does anyone think that this tool is well-tuned?*
  - FNAL support for making rooms/equipment useable is not good.
    - *Not Sheila, she's great.*
    - *I think that FNAL-based people still don't appreciate the frustration of the outsider*
- *CERN document server: Frank Filthaut has brought up*
  - but some DØ groups will not use it.

## tool development is complicated

Groupware developers face new challenges,  
e.g.

*for single-user applications—as in baseball—if 3/10 consumers adopt software, it's a resounding success. But, for groupware, such an adoption rate is a disaster.*

- Critical mass acceptance is a make-or break issue: everyone must use.

- Early adopters face a Prisoner's Dilemma

- *“Media Stickiness” is the phenomenon of smart people resisting adoption of new media tools*

- *Developers must “live” with early adopters, or they will quit*

*who wins is a consideration...failed case in point: voice annotation tools*

- the speaker benefits, the subsequent user is hugely inconvenienced

Like I say, this subject is intensely studied  
*outside of High Energy Physics*

# should we be considering new tools?

The only historically successful groupware tool: email.  
Is that good, or is that bad?

*In the recent DØ ITR proposal among roughly 10 collaborators:*

- from 1/30/03-2/10/03, I received 1242 email messages related to proposal  
*many with an attached MS Word file.*
- on 2/10/03 alone, I received 188 email messages

*This is good?*

Does email serve all needs, equally well?

*There is evidence that text-based chat - you know, IM, like your kids use  
- is enormously productive among like-minded collaborators*

- One R&D product from LBNL DCP is java-based, secure, text chat  
*this is an example of an “awareness” tool - always on, among “buddies” - like  
hallway-encounter awareness*  
*could we try this: add greek, math symbol notation, hot linking, attachments?*  
*LBNL developers are eager for users and feedback*

*Web logs - “blogs” can be a source of focused, thematic interaction  
among collaborators*

- a permanent, single home for a document + chronological or threaded discussions  
of the document

# people/organization challenges

## Two kinds of challenges:

### 1. *Physics group organization*

- currently “blind” to explicit on-site/off-site locations
- can Physics/ID group organization make use of Regions?
- I think that radical organizational change is required to efficiently mobilize the entire collaboration over our 4 continents

### 2. *Collaborative tools*

- we cannot wait for all sexy collaborative alternatives
- what few tools might improve our communications and efficiency?

#### **Beaune:**

1. What is the optimum Physics Group organization to include off-site people? Can this bootstrap from the RAC concept?
2. Should we consider collaborative tools integration into DØ?  
*chat? blogs? document server? enhanced VNC? electronic logbooks? white boards?*
3. Will people use them? Or, is everyone happy with email.



## conclusions

**DØ faces qualitatively new challenges to success in analyzing Run II**

- challenges which are technical - RAC's, SAM-Grid are steps*
- challenges which are brand new: organizing hundreds of people*

**The physics is worth the effort**

*but, we need to try new things*

*we need to evolve our old ways of analyzing experiments*

*we need to all adopt agreed upon tools*