

# ***SAGrid 2.0: History, Lessons, Evolution***

Bruce Becker | [bbecker@csir.co.za](mailto:bbecker@csir.co.za)  
Coordinator, SAGrid  
Meraka Institute, CSIR



# Outline


- This talk... apologies to Magritte, Pope Francis, Octocat, and you.
- The grid -
  - What is it ? What do people think it is ? What do people think it does ? What does it really do ?
- Failures, Bottlenecks
- Technology changes
- Methodology changes
  - Site deployment
  - Application porting



# Ceci n'est pas une presentation



*ceci n'est pas une chaussure*



*Ceci n'est pas une MOUSTACHE*



*Ceci n'est pas ceci*



*Ceci n'est pas un bureau*



*Ceci n'est pas un vecteur.*



*Ceci n'est pas  
une carte  
d'anniversaire*




*Ceci n'est pas une nike*



*Ceci n'est pas un neurone.*



*Ceci n'est pas la guerre.*



*ceci n'est pas une pipe.*



*Ceci n'est pas une couch gag.*




*Ceci n'est pas une pipe.*



*Ceci n'est pas un nuage.*



*Ceci n'est pas une browser.*



*Ceci n'est pas une pipe.*



*Ceci n'est pas Magritte*



*Ceci n'est pas un chapeau.*



# Perceptions

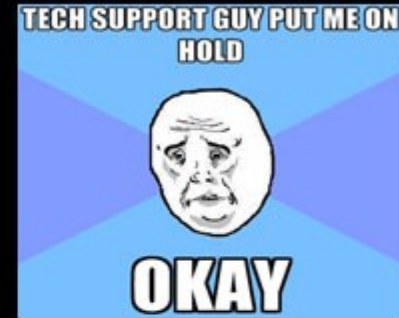
## IT



What my friends think I do



What my mom thinks I do



What society thinks I do



What my boss thinks I do



What I think I do



What I really do



# What is the grid





Ok... can it make coffee ?

# What is "SAGrid" ? (2007)



## **"The" :**

- there is one major effort to deploy computing resources, in a coherent and well-managed way.
- This does not exclude smaller or less-significant "research" projects, which may be integrated later or interoperate with the national infrastructure

## **"South African National" :**

- The infrastructure is hosted and maintained within the borders of South Africa and the operations team is South African
- Expertise is locally built, exploiting international contacts in operations and middleware
- A single operations team (a Virtual Organisation) is responsible for the operational state of the entire grid.

## **"Grid"**

- This is a federation of computational resources, respecting the definition of a compute grid and a data grid
- The computing resources are and remain the property of the site which paid for them, although they are shared amongst all users
- The kind and level of resource sharing is agreed to on a virtual organisation basis



# Goal

***to enable anyone  
with the right credentials anywhere in South Africa,  
to access the  
entire national and, if it is the case, international  
cyberinfrastructure and software  
in a way that is uniform, coherent and easy-to-use  
and to provide a  
federated, flexible and dynamic  
tool for collaboration and research.***





# Goal 2.0

*to enable anyone  
with the right credentials anywhere in ~~South Africa~~  
Africa and Arabia*

*to access the  
entire national and, if it is the case, international*


*Cyberinfrastructure (data, compute, software,  
metadata, support, etc services)  
in a way that is uniform, coherent and easy-to-use*

*and to provide a  
federated, flexible and dynamic*

*tool for collaboration and research.`*



# What we wanted to do





# What we wanted to do

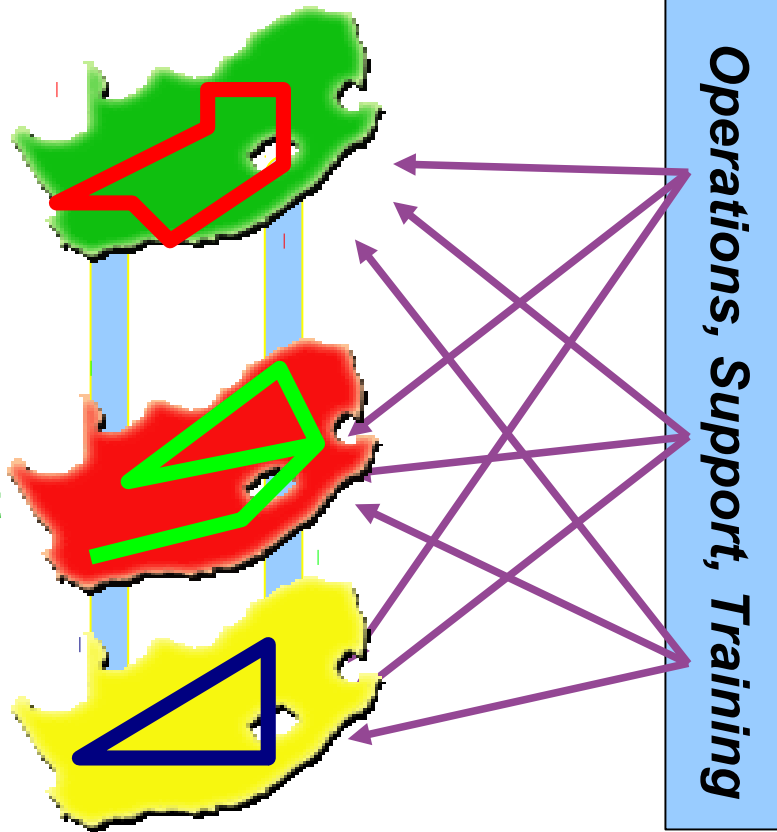
- IT Departments of Universities centralise computing/data resources on site
  - Owners of equipment form technical collaboration as the basis of the grid – Joint Research Unit Agreement (JRU)
  - Activities coordinated by neutral 3rd party (Meraka)
  - Strategy and priorities driven from within the JRU, based on institutional needs
- Deploy standards-based interface to each resources using widely-used middleware
  - Train IT staff to operate it
  - Form national operations team
- Create catch-all VO and sign up users
- Identify and port applications
- ...
- Science !



# What we wanted to do

**National  
Research  
Network  
connects**

**HPC resources  
and  
scientific equipment**  
which  
**self-organising  
user communities**  
use to conduct  
their research



Users access  
instruments,  
software and  
computing  
resources  
independent  
of their  
location in a  
self-organised  
way,  
transparently  
across  
infrastructures

e-infrastructure







# SAGrid activities

- What does SAGrid do ? **”HPC and data storage + other esoteric ”grid” stuff” right ?**
- Nope – we tried to make a full platform for collaboration
  - Virtual Organisation service – <https://voms.sagrid.ac.za>
  - Tracking system – <https://ops.sagrid.ac.za>
    - Projects, tasks
    - Help and feedback
  - Agenda service – <https://indico.sagrid.ac.za>
  - Digital repository – <https://docs.sagrid.ac.za>
  - Etc..
- Were installed to help coordinate the operations of the infrastructure, but could be used by all user communities

# How it was (not) supposed to be managed





# How it was (not) supposed to be managed

But neither  
some  
mysterious,  
central  
dictatorship



*Ceci n'est pas une pape.*





# How it was supposed to be managed

- JRU was supposed to be a decision-making body, allowing **collective** action
  - Decisions are made by JRU, implemented bottom-up by members
  - Coordinated by Meraka
- In the case of participation to collaborative research or development projects, CSIR Meraka could represent the collaboration
- Institutes maintain full ownership and control of their equipment, but agree to share it



# How it's really managed

- JRU takes no proactive action
- Members act individually, rather than collectively.
  - Incoherent hardware acquisition – seen as waste or duplication by NRF
  - Inconsistent (inexistent) communications strategy – seen as institutional only
- Decisions are made by coordinator, implemented “brute-force” top-down
  - Slow, meets with resistance due to conflicting priorities
- No collective participation to e-Infrastructure support actions



# What we really do

- Sites differ significantly in capacity and resources
- Separated resources at any given site require specific attention – does not scale
- Loosely-coupled experts form national operations team – initial barrier to entry for new sites remains high
- Inconsistent monitoring and accounting – no accurate view of usage and performance



# How users were supposed to use it



# A scenario : this scientist walks into a bar...






# What users want

Some users want highly varied, modular application selection

Amazing infrastructure



Highly trained support



Vertically integrated  
Highly specialised applications





# What users get sometimes









# What went wrong ?

I'VE LEARNED SO MUCH  
FROM MY MISTAKES...

I'M THINKING OF MAKING  
A FEW MORE.



# Problems help you learn





# Failures : HPC vs grid

- There is a widely-held belief that HPC and grid computing are mortal enemies
  - HPC = big machine for supercomputing / MPI
  - GRID = lots of little machines for "HEP"
- This is highly counterproductive in SA
  - We don't have any "supercomputers" - all cluster-based hybrid architectures
  - Advances in middleware, schedulers and workload management systems have made this distinction irrelevant
  - Virtual environments make this distinction even more useless
  - Feeling that having to support "grid" and "HPC" is double effort – strain on site, have to choose one.



# HPC vs grid – for the record

- The grid is a federation of resources and services, one kind of which is HPC
- Middleware is a "job transport mechanism"
- Why did this not work in SA ?
  - Proper advertising of resources
  - Common environment for applications
  - No effort on standardising priority and scheduling policies
  - Very little sharing of user-gridops interaction to learn best practices



# Failures - Site X

- Very good case to join the infrastructure
  - Well-defined set of applications, widely used
  - Compact team using the site
  - Site admin initially trained at EPIKH school
- Looked good for a while...
  - Site deployed, tested, ready
  - Hardware upgrade + middleware upgrade overloaded system admins
  - Site was removed from the infrastructure



# Why ?

- GridOps was seen as an "extra" activity
- "price" you pay to join the grid – but what are you getting in return ?
- No help was called in for problems
- Didn't work enough on having relevant applications available to users across the grid
- Unbalanced equation : using local resources, but supporting remote users – no thanks.



# Failures – SAAO, SANBI, tLABS

- National facilities with large pools of users
- Commonly-used applications
- Local expertise in HPC and user support
- Almost zero takeup...

# Why was it so hard to bring new users on board ?



- Culture shock
  - Fundamental misconception of what the grid can do / what it is
  - Missed opportunities of joining existing Virtual Organisation – user feels alone
  - No specific message – “the grid can do everything !” yes, but how do I do this specific thing
  - Static, irrelevant training material – doesn't speak to the user
  - One-way communication – no feedback channel → no improvement
  - Responsiveness too low
- Barriers to entry
  - x.509 certificates, new technology/methodology
  - Training model





# New User Bottlenecks

- I want to get started !
  - Wait for your certificate
- I want to get started.
  - Learn a new paradigm and a new language
- I want to get started...
  - Tell us about your code and we'll put it in the list of applications to be ported
- Bye...



# New Site Bottlenecks

- Q: What's this grid thing ?
  - A: Set up your cluster, rtfm, deploy middleware, give us your bdii and you're good.
- Q: Yes, but what *is* this grid thing ?
  - A: You have 12 tickets open, what's the problem ?
- Q: It's for physicists, right ? But I have astronomers asking for HPC...
  - A: Ok, we're going through a middleware upgrade – when are you going to be upgrading ?



# Parlez-vous grid ?

It's not you, it's ~~me~~

*...YOUR MESSAGING*



# Changes to the playing field

- New middleware stack
- Highly improved technology provider, but some disruption





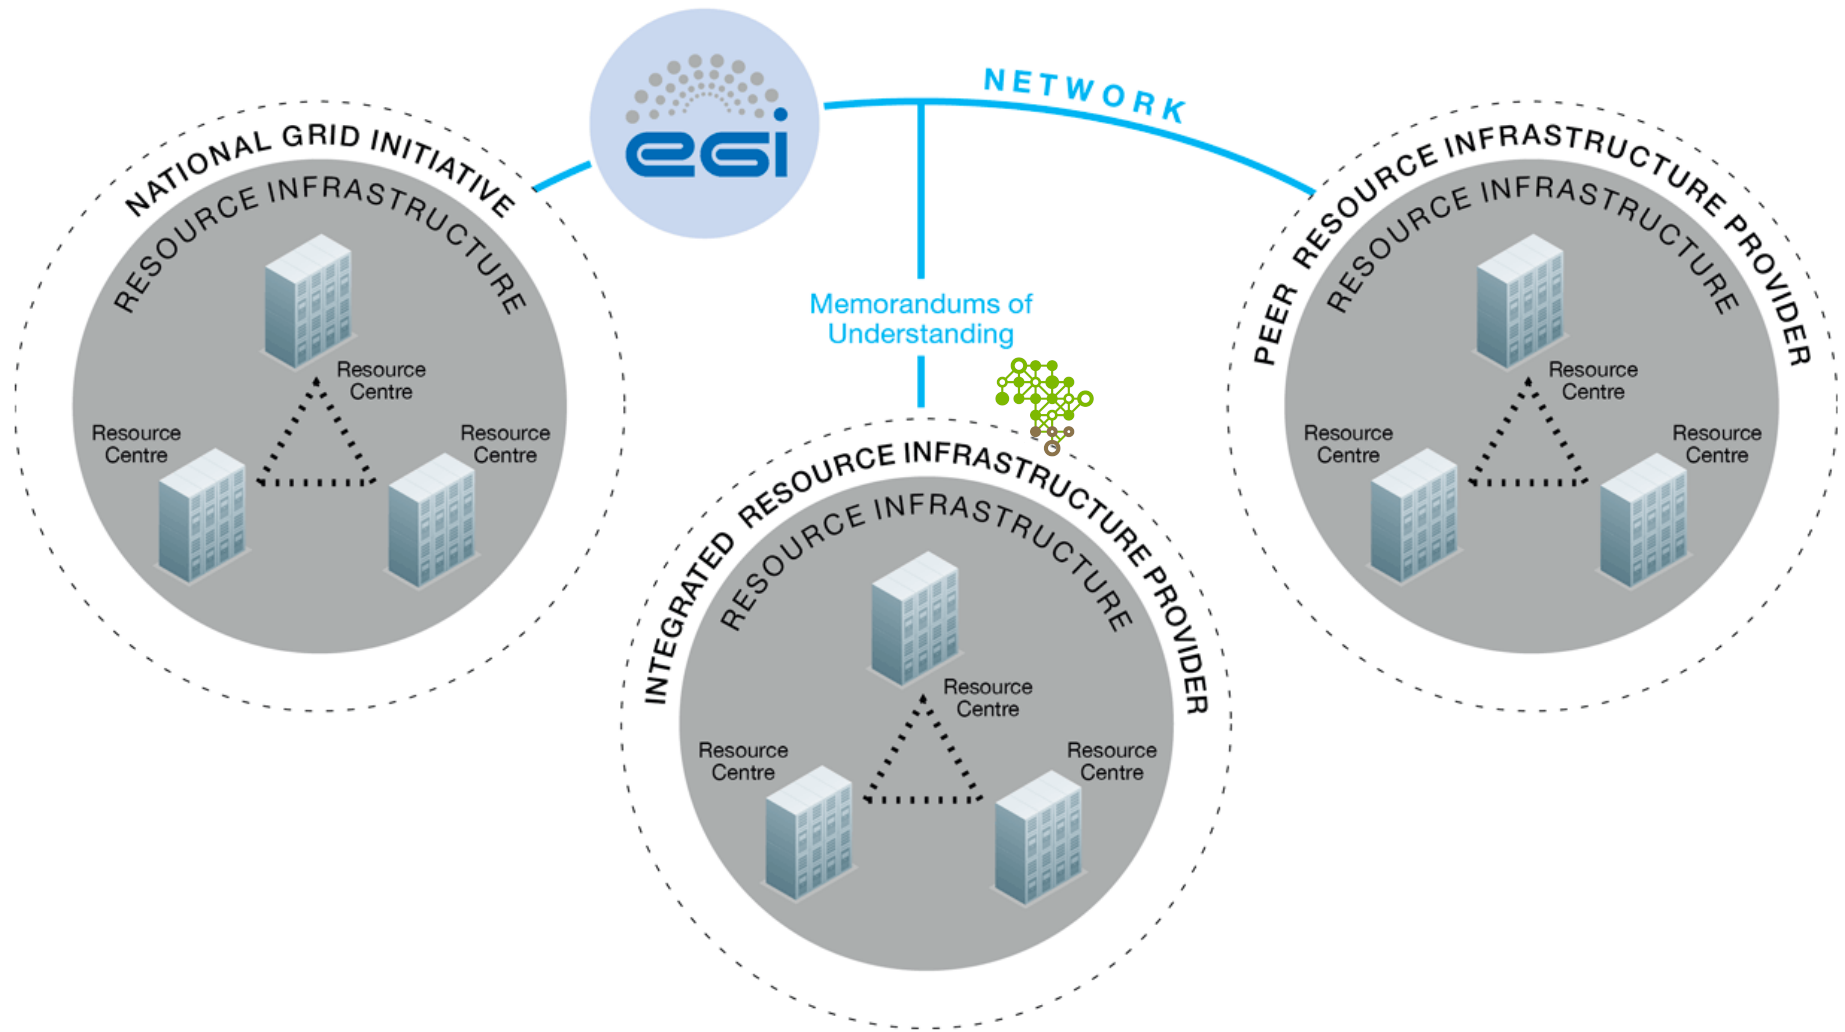
# Changes to the playing field

- More formal technical basis to the collaboration
  - OLA/SLA basis to operational infrastructure and services
- Better understanding of the state and nature of the services provided
- Easier to add, publish services and collaborate: e.g. :
  - Monitoring
  - Security
  - Web services
  - Identity Services





# An interoperable infrastructure





# NGI\_ZA



What is an NGI?

### Contacts



E-Mail	sagrid-ops@lists.sagrid.ac.za
ROD E-Mail	
Helpdesk E-Mail	
Security E-Mail	sagrid-site-security@lists.sagrid.ac.za

### Project memberships



EGI
-----

### Scope(s)

EGI	?
-----	---

### 6 Sites



Name	Certification Status	Production Status	Scope(s)
ZA-UJ	Certified	Production	EGI
ZA-CHPC	Certified	Production	Local
ZA-UFS	Certified	Production	Local
ZA-MERAKA	Certified	Production	EGI
ZA-UCT-ICTS	Certified	Production	Local
ZA-WITS-CORE	Certified	Production	EGI

### 8 Users



Name	Role
Bruce Becker	NGI Operations Manager
Ntuthuko Sambo	Regional Staff (ROD)
Christopher Lee	Regional Staff (ROD)
Roderick Meei	NGI Security Officer





# Who is the grid ?







# New Services - SA CA

https://security.sanren.ac.za/CA/



## South African National Grid Certification Authority

### Home

### South African National Grid CA certificate

### Request a Personal certificate

### Request a Robot certificate

### Request an Host certificate

### Renew Personal or Robot certificate

### Check certificates

### Certificate Revocation List

The **Certificate Authority** issues one year Personal and Robot public key certificates (compliant with the X.509 standard).

For any problems, please contact:

e-mail: [bbecker@csir.co.za](mailto:bbecker@csir.co.za)


In order to view the South African National Grid CA public certificate and/or save it in your web browser (which is necessary to validate your personal certificate) click on **South African National Grid CA certificate** in the left part of this page.

Please, note that Mozilla, Chrome and Internet Explorer are the only presently supported web browsers. The use of any other web browsers could induce some visualization mismatches and/or server misbehaviours and is not currently suggested.

## Joint SAGrid/SANREN service for all SA e-Science



# New Services - SA CA Accreditation



## Meetings

(meetings from our sister PMAs are available at: [the TAGPMA conference site](#) and [the APGridPMA F2F meeting site](#))

Date	Location	Details
January 2014	Abingdon, Oxfordshire, UK hosted by RAL STEC	Monday 13 - Wed 15 January 2014, with SCI meeting on 15th and 16th
September 2013	Bucharest, Romania hosted by ROSA	Monday 9 - Wed 11 September 2013
May 2013	Kyiv, Ukraine hosted by UGrid	Monday 13 - Wed 15 May 2013





# Easier Access to services

<https://idp.sagrid.ac.za>

Home

Service Providers

About

Register

Login

**SOUTH  
AFRICAN  
NATIONAL  
GRID**



## Home

The South African Catch-All Identity Provider is a SAML standard-based service powered by Shibboleth and jointly managed by SAGrid and SANREN that allows people to use a single identity when they sign into web pages and Science Gateways that require user authentication.

This IdP is a member of Grid IDentity Pool (GrIDP) federation and currently authenticates the users of several Service Providers

If your institute does not have an Identity Provider which forms part of a national or catch-all federation, you can register to the catch-all IdP

If you are interested in either using this IdP to authenticate users of your website(s), or registering it as an Identity Provider of your Identity Federation(s), please contact Bruce Becker at [bbecker\\_at\\_csir.co.za](mailto:bbecker_at_csir.co.za).

For more information, please see the ei4Africa project website.

Brought to you by



***Talk to Siju Mammen***  
***smammen@csir.co.za***





# Science Gateways

The screenshot shows the Africa Grid Science Gateway website. At the top, there is a navigation bar with a logo of Africa on the left, the title "Africa Grid Science Gateway" in the center, and logos for Africa, CTSG, and the European Union on the right. Below the navigation bar is a menu with "Welcome" and "Applications". The main content area features a banner for "STATISTICAL COMPUTING AND GRAPHICS ON DISTRIBUTED ENVIRONMENTS" with the R logo. Below the banner are two tabs: "Welcome" and "Contributors". The "Welcome" tab is active, showing text about the gateway and access. The "Contributors" tab is also visible. At the bottom, there is a red text overlay with contact information and a URL.

Register Sign In

Welcome Applications

STATISTICAL COMPUTING AND GRAPHICS ON DISTRIBUTED ENVIRONMENTS

Welcome Contributors

The Africa Grid Science Gateway provides access to various scientific applications identified by the e-Infrastructure Consortium for Africa (ICA). The access to the Africa Science Gateway organisation you belong to has the page. Otherwise, you can go to the [GridDP](#) "catch-all" federation. On the right corner of the page

**Talk to Albert van Eck ([vanecka@ufs.ac.za](mailto:vanecka@ufs.ac.za))**  
**Join CTSG [sg-license@ct.infn.it](mailto:sg-license@ct.infn.it)**  
**<http://www.catania-science-gateways.it/>**

individuals wanting to contribute to the



# Science Gateways

Register Sign In

Welcome Applications

STATISTICAL COMPUTING AND GRAPHICS ON DISTRIBUTED ENVIRONME

SOUTH AFRICAN NATIONAL GRID

Log in to Africa Grid Science Gateway

BruceBecker

.....

Login

- > Forgot your password?
- > Need Help?

## Welcome

The Africa Grid Science Gateway is a standard-based web 2.0 demonstr applications identified by the [e4Africa project](#) and execute them on a wc

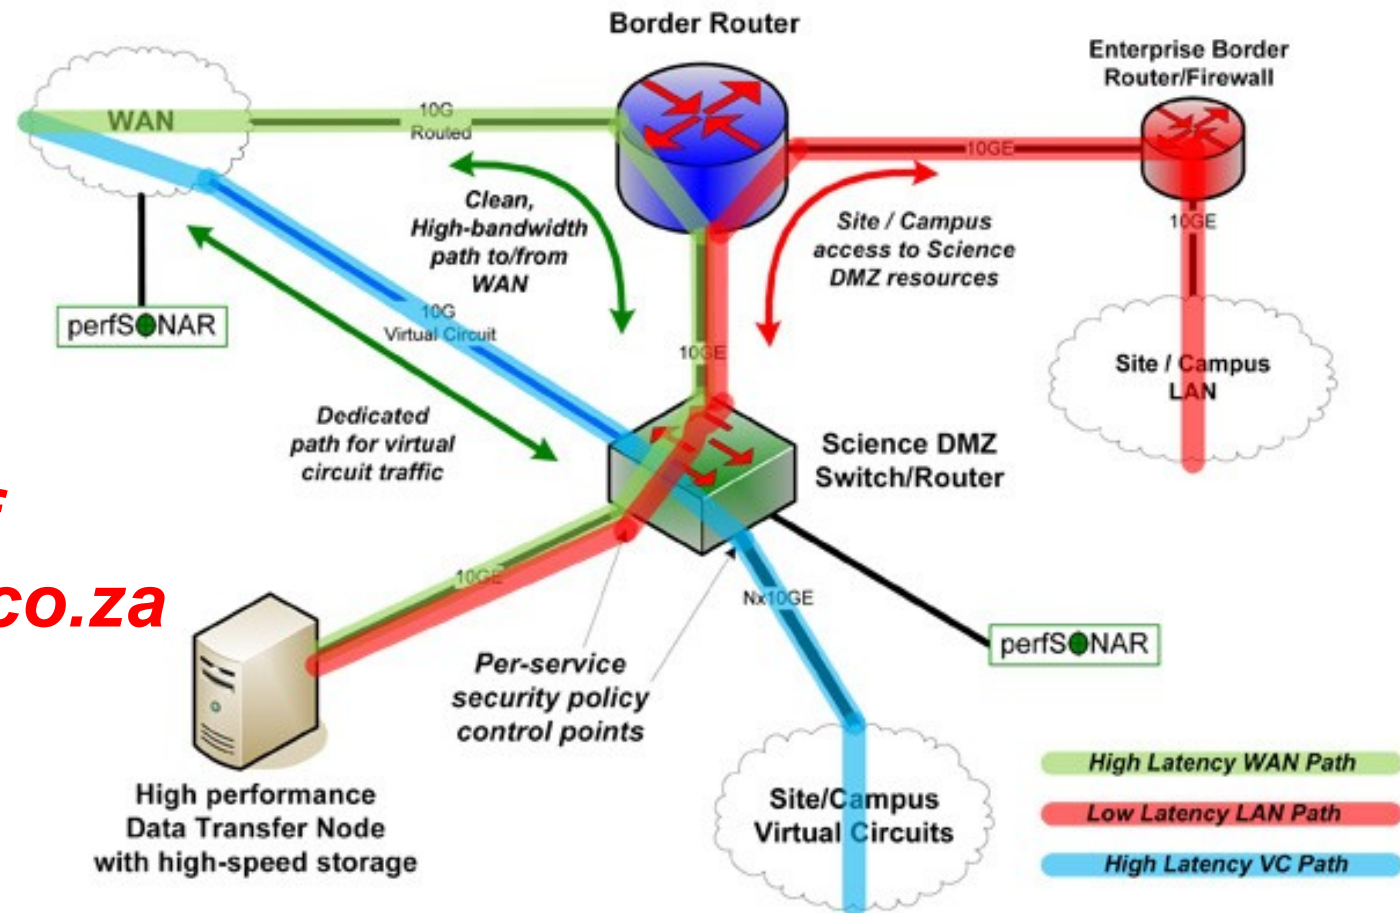
The access to the Africa Science Gateway requires federated credentials organisation you belong to has an Identity Provider, click on the "Sign In" the page. Otherwise, you can get federated credentials registering to the [GrIDP](#) "catch-all" federation. Once you have gotten your credentials, click right corner of the page

In order to run an application, select it from the Applications menu above. New applications can also be proposed to be



# Data: Movement

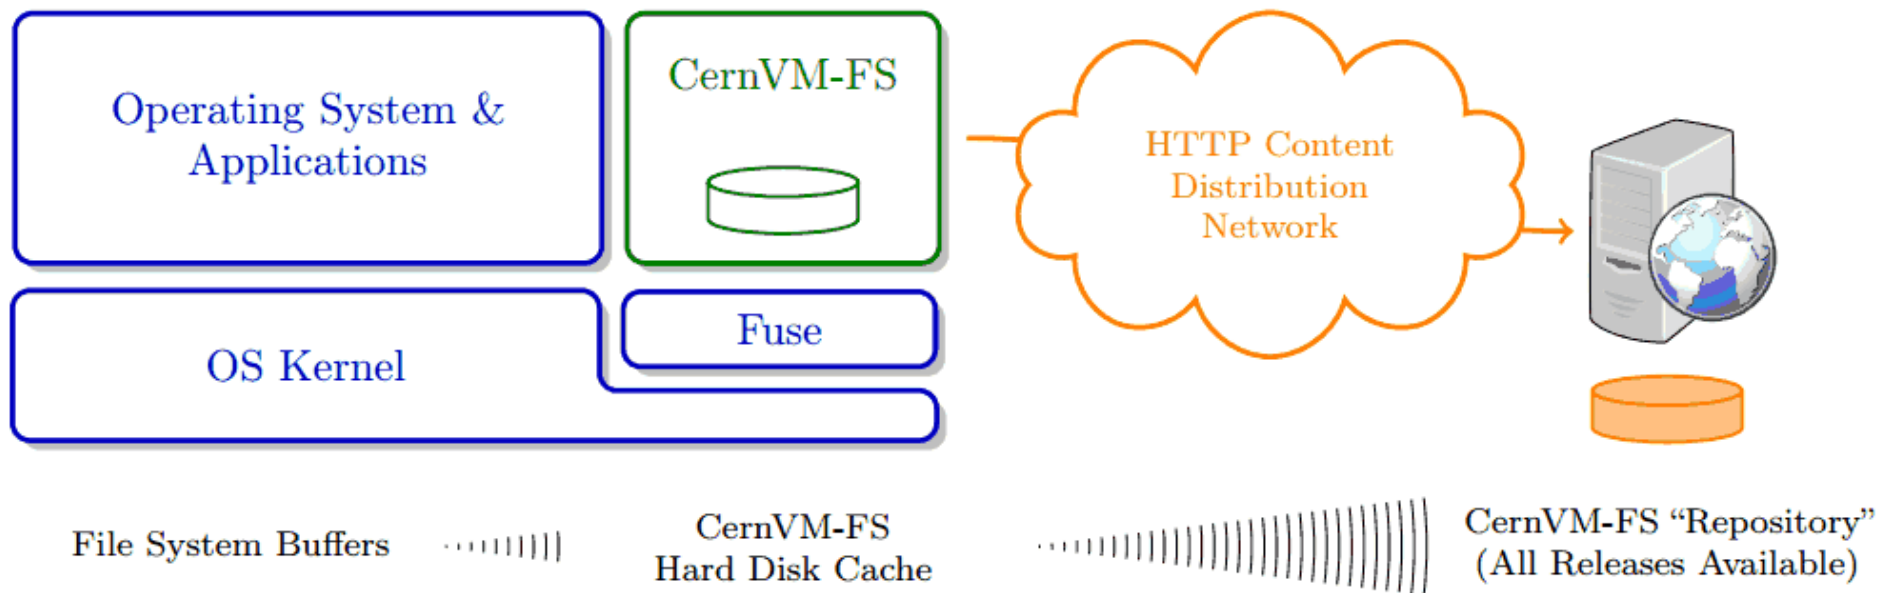
- Why aren't we moving data at the speed of light?



**Talk to  
Simeon Miteff  
smiteff@csir.co.za**



# Application distribution : CERNVMFS





# Make it easier to contribute

- I have something I want on the grid ...
- I have some code that the ops team can use...
- etc

SOUTH AFRICAN NATIONAL GRID

logged in as /C=IT/O=INFN/OU=Personal Certificate/L=ZA-MERAKA/CN=Bruce Becker | [Logout](#) | [Preferences](#) | [Help/Guide](#) | [About Trac](#)

Documentation | Timeline | Roadmap | **Browse Source** | View Tickets | Open a new ticket | Search | Calendar | Admin | Gantt chart

[Subversion Location](#)


### Repository Index

Name ▲	Size	Rev	Age	Author	Last Change
▼ <b>repo</b>		544	5 months	anonymous	Added 'Makefile' to tarball. Excluded '.svn' files from tarball ...
▶ <b>branches</b>		2	3 years	root	initial import
▶ <b>prod</b>		402	15 months	anonymous	added a new branch for production code
▶ <b>tags</b>		38	3 years	anonymous	tagging for the potch tutorial
▶ <b>trunk</b>		540	6 months	anonymous	
▶ <b>users</b>		544	5 months	anonymous	Added 'Makefile' to tarball. Excluded '.svn' files from tarball ...
▶ <b>software_installation</b>			-	-	TracError: GIT version 1.5.4.7 installed not compatible(need >= 1.5.6)





# Recognise individuals... wider trust base



**Bruce Becker**  
brucellino

Meraka Institute, CSIR  
Pretoria, South Africa  
 [bbecker@csir.co.za](mailto:bbecker@csir.co.za)  
 <http://www.sagrid.ac.za>  
 Joined on Aug 08, 2012

**1** follower  
**7** starred  
**8** following

[Contributions](#) [Repositories](#) [Public Activity](#) [Edit Your Profile](#)

### Popular repositories

- [SAGridIdP](#) Code for the Catch-All IdP of South Africa 1 ★
- [SAGridDemo](#) A repository of demo applications to sho... 1 ★
- [SAGridApps](#) SAGrid Applications 0 ★
- [Quarks-Matter](#) Quarks Matter - the story of how and why. 0 ★
- [test](#) playground 0 ★

### Repositories contributed to

- [SAGridOps/SoftwareInstallation](#) The installation scripts for SAGrid Applic... 1 ★
- [SAGridOps/ServiceDocumenta...](#) A repository to contain the documentati... 0 ★
- [SAGridOps/Documents](#) A place for public documents 0 ★

### Your Contributions

Summary of Pull Requests, issues opened and commits. [Learn more.](#) Less More



# Decentralise the team



## SAGrid Operations team

SAGridOps

South Africa  
sagrid-ops@lists.sagrid.ac.za  
<http://www.sagrid.ac.za>  
Joined on Jul 05, 2013

6 public repos   0 private repos   1 member

Repositories   Members

Edit SAGridOps's Profile

Find a repository...

Search

All Public Private Sources Forks Mirrors



### grid-services-deployment

forked from [auth-scc/grid-services-deployment](#)

Example ansible playbooks that are used to deploy Grid services

Last updated 12 days ago

Shell   ★ 0   1



### Documents

A place for public documents

Last updated 16 days ago

★ 0   0



### DevOps

Top-level directory of devops scripts for SAGrid

Last updated 20 days ago

★ 0   0



### SoftwareInstallation

The installation scripts for SAGrid Application deployment.

Last updated 2 months ago

Shell   ★ 1   0



### ServiceDocumentation



A repository to contain the documentation and planning for various services. Includes technical documentation, prototypes, etc.

★ 0   0





# Collaborate with code




Repositories Members

Edit SAGridOps's Profile

Find a repository... Search

All Public Private Sources Forks Mirrors

 **grid-services-deployment**  
 forked from [auth-scc/grid-services-deployment](#)  
 Example ansible playbooks that are used to deploy Grid services  
 Last updated 12 days ago

Shell ★ 0 1

**SAGrid Operations team**  
SAGridOps


 **Documents** ★ 0 0  
 A place for public documents  
 Last updated 16 days ago

South Africa  
 sagrid-ops@lists.sagrid.ac.za  
<http://www.sagrid.ac.za>  
 Joined on Jul 05, 2013


 **DevOps** ★ 0 0  
 Top-level directory of devops scripts for SAGrid  
 Last updated 20 days ago

6 public repos 0 private repos 1 member

 **SoftwareInstallation** Shell ★ 1 0  
 The installation scripts for SAGrid Application deployment.  
 Last updated 2 months ago

 **ServiceDocumentation** ★ 0 0  
 A repository to contain the documentation and planning for various services. Includes technical documentation, prototypes, etc.





# Let the robots do the work

- Define what we want to deploy – let the experts take care of how to deploy
- **DevOps** paradigm – same review/tag/release mechanisms on operations code as we have for scientific applications
  - Teach a marketable skill
  - Allow specialisation
  - Enable remote management of complex services
  - Ensure that published methodology is adopted methodology





# Quality Control and feedback



- Ensure that requested applications are included in the repo
- Provide testing and QA infrastructure
- Self-serve to users

***Talk to Fanie Riekert***  
**fanie.riekert@gmail.com**



# Well done, you made it to the end of the talk

- "SAGrid" is now a misnomer – have you got a suggestion ?
  - NGI\_ZA + lots of other services in a federation
  - A lot of assumptions about **what** and **how** need to be updated
- Underlying infrastructure and core services need to be reliable – measured and improved
- Peer MoU with EGI.eu
  - greatly increased reliability/availability and scope of services
  - Direct collaboration with other NGI/NRENS
- Eat your own dog food – if it tastes bad, don't be afraid to change it
- Remove barriers to access, usage and development
- Predicting, reporting, managing, implementing infrastructure services



# Thank you !

- It's a community... there's no "us" and "you"
  - SAGrid on Facebook and Twitter
    - <https://www.facebook.com/SAGrid>
    - <https://www.twitter.com/TheSAGrid>
- There's place for everyone and there doesn't necessarily have to be only 1 leader





# ***SAGrid 2.0: History, Lessons, Evolution***

Bruce Becker | [bbecker@csir.co.za](mailto:bbecker@csir.co.za)  
Coordinator, SAGrid  
Meraka Institute, CSIR





## Outline

- This talk... apologies to Magritte, Pope Francis, Octocat, and you.
- The grid -
  - What is it ? What do people think it is ? What do people think it does ? What does it really do ?
- Failures, Bottlenecks
- Technology changes
- Methodology changes
  - Site deployment
  - Application porting



Bruce Becker: Coordinator, SAGrid | [bbecker@csir.co.za](mailto:bbecker@csir.co.za) | <http://www.sagrid.ac.za>

This is the outline of my presentation. I will describe some of the philosophy, interventions and lessons we have learned in South Africa regarding the development of an integrated e-Infrastructure; then I will briefly give a snapshot of the status of such activities in Africa.

Thereafter, I will go into some consideration on how to sustainably support and grow these, leading into the core of my presentation – the Africa and Arabia Regional Operations Centre. Finally, a summary and outlook will be given.



# Ceci n'est pas une presentation



*ceci n'est pas une chaussure*



*Ceci n'est pas une lunettes*



*Ceci n'est pas ceci*



*Ceci n'est pas un bureau*



*Ceci n'est pas un vecteur.*



*Ceci n'est pas une carte d'anniversaire*




*Ceci n'est pas une nike*



*Ceci n'est pas un neurone.*




*Ceci n'est pas la guerre.*



*ceci n'est pas une pipe.*



*Ceci n'est pas une couch gag*



*Ceci n'est pas une pipe.*



*Ceci n'est pas un nuage.*



*Ceci n'est pas une browser.*



*Ceci n'est pas une pipe.*



*Ceci n'est pas Magritte*



*Ceci n'est pas un chapeau.*



# Perceptions


## IT



What my friends think I do



What my mom thinks I do



What society thinks I do



What my boss thinks I do



What I think I do



What I really do




# What is the grid



Ok... can it make coffee ?







# Goal

*to enable anyone  
with the right credentials anywhere in South Africa,  
to access the  
entire national and, if it is the case, international  
cyberinfrastructure and software  
in a way that is uniform, coherent and easy-to-use  
and to provide a  
federated, flexible and dynamic  
tool for collaboration and research.*



## Goal 2.0

*to enable anyone  
with the right credentials anywhere in ~~South Africa~~  
Africa and Arabia*

*to access the  
entire national and, if it is the case, international*

*Cyberinfrastructure (data, compute, software,  
metadata, support, etc services)  
in a way that is uniform, coherent and easy-to-use*

*and to provide a  
federated, flexible and dynamic*

*tool for collaboration and research.`*







What we wanted to do



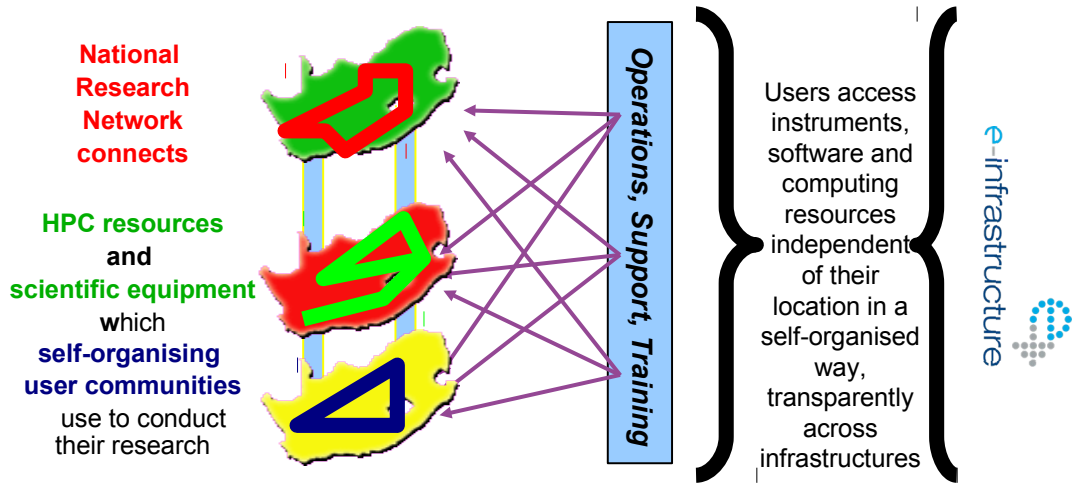


# What we wanted to do

- IT Departments of Universities centralise computing/data resources on site
  - Owners of equipment form technical collaboration as the basis of the grid – Joint Research Unit Agreement (JRU)
  - Activities coordinated by neutral 3rd party (Meraka)
  - Strategy and priorities driven from within the JRU, based on institutional needs
- Deploy standards-based interface to each resources using widely-used middleware
  - Train IT staff to operate it
  - Form national operations team
- Create catch-all VO and sign up users
- Identify and port applications
- ...
- Science !



# What we wanted to do








# SAGrid activities

- What does SAGrid do ? **”HPC and data storage + other esoteric ”grid” stuff” right ?**
- Nope – we tried to make a full platform for collaboration
  - Virtual Organisation service – <https://voms.sagrid.ac.za>
  - Tracking system – <https://ops.sagrid.ac.za>
    - Projects, tasks
    - Help and feedback
  - Agenda service – <https://indico.sagrid.ac.za>
  - Digital repository – <https://docs.sagrid.ac.za>
  - Etc..
- Were installed to help coordinate the operations of the infrastructure, but could be used by all user communities

# How it was (not) supposed to be managed



# How it was (not) supposed to be managed



But neither  
some  
mysterious,  
central  
dictatorship



*Ceci n'est pas une pape.*



## How it was supposed to be managed

- JRU was supposed to be a decision-making body, allowing **collective** action
  - Decisions are made by JRU, implemented bottom-up by members
  - Coordinated by Meraka
- In the case of participation to collaborative research or development projects, CSIR Meraka could represent the collaboration
- Institutes maintain full ownership and control of their equipment, but agree to share it





# How it's really managed

- JRU takes no proactive action
- Members act individually, rather than collectively.
  - Incoherent hardware acquisition – seen as waste or duplication by NRF
  - Inconsistent (inexistent) communications strategy – seen as institutional only
- Decisions are made by coordinator, implemented "brute-force" top-down
  - Slow, meets with resistance due to conflicting priorities
- No collective participation to e-Infrastructure support actions



## What we really do

- Sites differ significantly in capacity and resources
- Separated resources at any given site require specific attention – does not scale
- Loosely-coupled experts form national operations team – initial barrier to entry for new sites remains high
- Inconsistent monitoring and accounting – no accurate view of usage and performance



# How users were supposed to use it

A scenario : this scientist walks  
into a bar...






# What users want

Some users want highly varied, modular application selection

Amazing infrastructure



Highly trained support



Vertically integrated  
Highly specialised applications





# What users get sometimes







# What went wrong ?

I'VE LEARNED SO MUCH  
FROM MY MISTAKES...

I'M THINKING OF MAKING  
A FEW MORE.





# Problems help you learn





## Failures : HPC vs grid

- There is a widely-held belief that HPC and grid computing are mortal enemies
  - HPC = big machine for supercomputing / MPI
  - GRID = lots of little machines for "HEP"
- This is highly counterproductive in SA
  - We don't have any "supercomputers" - all cluster-based hybrid architectures
  - Advances in middleware, schedulers and workload management systems have made this distinction irrelevant
  - Virtual environments make this distinction even more useless
  - Feeling that having to support "grid" and "HPC" is double effort – strain on site, have to choose one.



## HPC vs grid - for the record

- The grid is a federation of resources and services, one kind of which is HPC
- Middleware is a "job transport mechanism"
- Why did this not work in SA ?
  - Proper advertising of resources
  - Common environment for applications
  - No effort on standardising priority and scheduling policies
  - Very little sharing of user-gridops interaction to learn best practices



## Failures - Site X

- Very good case to join the infrastructure
  - Well-defined set of applications, widely used
  - Compact team using the site
  - Site admin initially trained at EPIKH school
- Looked good for a while...
  - Site deployed, tested, ready
  - Hardware upgrade + middleware upgrade overloaded system admins
  - Site was removed from the infrastructure



## Why ?

- GridOps was seen as an "extra" activity
- "price" you pay to join the grid – but what are you getting in return ?
- No help was called in for problems
- Didn't work enough on having relevant applications available to users across the grid
- Unbalanced equation : using local resources, but supporting remote users – no thanks.



## Failures – SAAO, SANBI, tLABS

- National facilities with large pools of users
- Commonly-used applications
- Local expertise in HPC and user support
- Almost zero takeup...

# Why was it so hard to bring new users on board ?



- Culture shock
  - Fundamental misconception of what the grid can do / what it is
  - Missed opportunities of joining existing Virtual Organisation – user feels alone
  - No specific message – “the grid can do everything !” yes, but how do I do this specific thing
  - Static, irrelevant training material – doesn't speak to the user
  - One-way communication – no feedback channel → no improvement
  - Responsiveness too low
- Barriers to entry
  - x.509 certificates, new technology/methodology
  - Training model



# New User Bottlenecks

- I want to get started !
  - Wait for your certificate
- I want to get started.
  - Learn a new paradigm and a new language
- I want to get started...
  - Tell us about your code and we'll put it in the list of applications to be ported
- Bye...





## New Site Bottlenecks

- Q: What's this grid thing ?
  - A: Set up your cluster, rfm, deploy middleware, give us your bdii and you're good.
- Q: Yes, but what *is* this grid thing ?
  - A: You have 12 tickets open, what's the problem ?
- Q: It's for physicists, right ? But I have astronomers asking for HPC...
  - A: Ok, we're going through a middleware upgrade – when are you going to be upgrading ?



Parlez-vous grid ?

**It's not you, it's ~~me~~**

*...YOUR MESSAGING*



# Changes to the playing field

- New middleware stack
- Highly improved technology provider, but some disruption





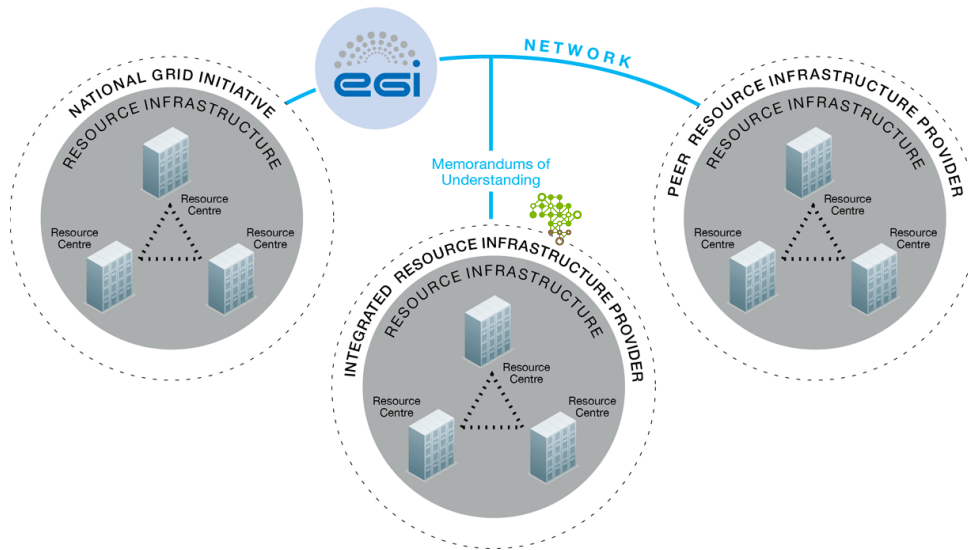
# Changes to the playing field

- More formal technical basis to the collaboration
  - OLA/SLA basis to operational infrastructure and services
- Better understanding of the state and nature of the services provided
- Easier to add, publish services and collaborate: e.g. :
  - Monitoring
  - Security
  - Web services
  - Identity Services





# An interoperable infrastructure





NGI\_ZA

Edit

What is an NGI?

Contacts

E-Mail	sagrid-ops@lists.sagrid.ac.za
ROD E-Mail	
Helpdesk E-Mail	
Security E-Mail	sagrid-site-security@lists.sagrid.ac.za

Project memberships

EGI
-----

Scope(s)

EGI ?
-------

6 Sites

Name	Certification Status	Production Status	Scope(s)
ZA-UJ	Certified	Production	EGI
ZA-CHPC	Certified	Production	Local
ZA-UFS	Certified	Production	Local
ZA-MERAKA	Certified	Production	EGI
ZA-UCT-ICTS	Certified	Production	Local
ZA-WITS-CORE	Certified	Production	EGI

8 Users

Name	Role
Bruce Becker	NGI Operations Manager
Ntuthuko Sambo	Regional Staff (ROD)
Christopher Lee	Regional Staff (ROD)
Dorothy Mool	NGI Security Officer



sagrid.ac.za



# Who is the grid ?





# New Services – SA CA

https://security.sanren.ac.za/CA/



## South African National Grid Certification Authority

[Home](#)

[South African National Grid CA certificate](#)

[Request a Personal certificate](#)

[Request a Robot certificate](#)

[Request an Host certificate](#)

[Renew Personal or Robot certificate](#)

[Check certificates](#)

[Certificate Revocation List](#)

The **Certificate Authority** issues one year Personal and Robot public key certificates (compliant with the X.509 standard).

For any problems, please contact:

e-mail: [bbecker@csir.co.za](mailto:bbecker@csir.co.za)

In order to view the South African National Grid CA public certificate and/or save it in your web browser (which is necessary to validate your personal certificate) click on **South African National Grid CA certificate** in the left part of this page.

Please, note that Mozilla, Chrome and Internet Explorer are the only presently supported web browsers. The use of any other web browsers could induce some visualization mismatches and/or server misbehaviours and is not currently suggested.

**Joint SAGrid/SANREN service  
for all SA e-Science**







# New Services – SA CA Accreditation



## Meetings

(meetings from our sister PMAs are available at: [the TAGPMA conference site](#) and [the APGridPMA F2F meeting site](#))

Date	Location	Details
January 2014	Abingdon, Oxfordshire, UK hosted by RAL-STFC	Monday 13 - Wed 15 January 2014, with SCI meeting on 15th and 16th
September 2013	Bucharest, Romania hosted by ROSA	Monday 9 - Wed 11 September 2013
May 2013	Kyiv, Ukraine hosted by UGrid	Monday 13 - Wed 15 May 2013





# Easier Access to services

<https://idp.sagrid.ac.za>

[Home](#)

[Service Providers](#)

[About](#)

[Register](#)

[Login](#)

**SOUTH  
AFRICAN  
NATIONAL  
GRID**



## Home

The South African Catch-All Identity Provider is a SAML standard-based service powered by Shibboleth and jointly managed by SAGrid and SANREN that allows people to use a single Identity when they sign into web pages and Science Gateways that require user authentication.

This IdP is a member of Grid Identity Pool (GridP) federation and currently authenticates the users of several Service Providers

If your institute does not have an Identity Provider which forms part of a national or catch-all federation, you can register to the catch-all IdP

If you are interested in either using this IdP to authenticate users of your website(s), or registering it as an Identity Provider of your Identity Federation(s), please contact Bruce Becker at [bbecker\\_at\\_csisr.co.za](mailto:bbecker_at_csisr.co.za).

For more information, please see the eI4Africa project website.

Brought to you by



**Talk to Siju Mammen**  
**[smammen@csir.co.za](mailto:smammen@csir.co.za)**





# Science Gateways

The Africa Grid Science Gateway: applications identified by the [GridPP](#) "catch-all" federation. On right corner of the page

**Talk to Albert van Eck ([vanecka@ufs.ac.za](mailto:vanecka@ufs.ac.za))**  
**Join CTSG [sg-license@ct.infn.it](mailto:sg-license@ct.infn.it)**  
**<http://www.catania-science-gateways.it/>**

individuals wanting to contribute to the

# Science Gateways



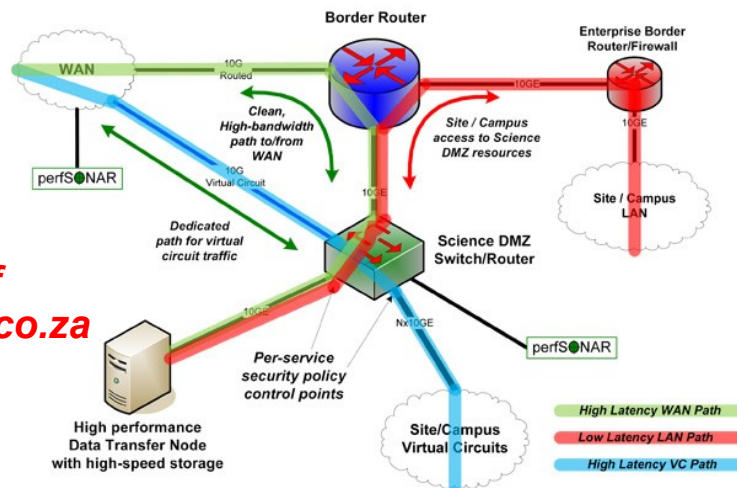
The screenshot shows the Africa Grid Science Gateway website. At the top, there is a navigation bar with a 'Welcome' tab selected and an 'Applications' tab. The main header features the text 'Africa Grid Science Gateway' in a large, green font, with a small globe icon to the left. Below the header, there is a banner for 'STATISTICAL COMPUTING AND GRAPHICS ON DISTRIBUTED ENVIRONMENTS' with a world map background. The central area contains a login form with the text 'Log in to Africa Grid Science Gateway'. The username field contains 'BruceBecker' and the password field is masked with dots. A red 'Login' button is positioned below the password field. To the right of the login form, there are links for '> Forgot your password?' and '> Need Help?'. The bottom left of the page has a 'Welcome' section with introductory text about the gateway's purpose and access requirements. The bottom right features logos for 'SANReN South African National Research Network' and 'SOUTH AFRICAN NATIONAL COMPUTE GRID'. A vertical sidebar on the right edge contains a 'Feedback' button.



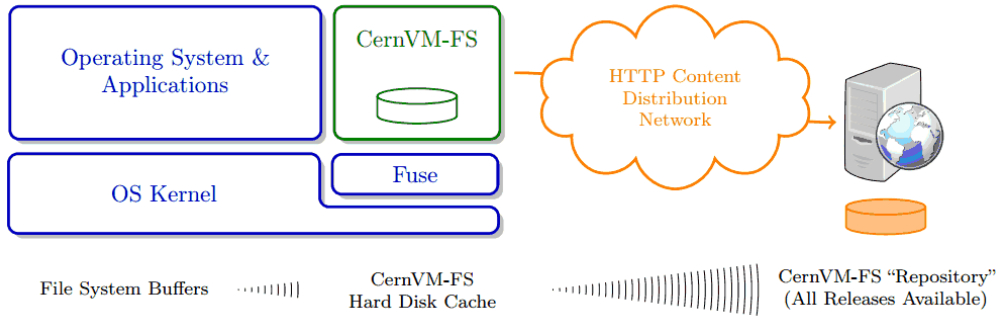
# Data: Movement

- Why aren't we moving data at the speed of light?

**Talk to  
Simeon Miteff  
smiteff@csir.co.za**



# Application distribution : CERNVMFS





# Make it easier to contribute

- I have something I want on the grid ...
- I have some code that the ops team can use...
- etc

SOUTH  
AFRICAN  
NATIONAL  
GRID



logged in as /C=IT/O=INFN/OU=Personal Certificate/L=ZA-MERAKA/CN=Bruce Becker | [Logout](#) | [Preferences](#) | [Help/Guide](#) | [About Trac](#)

[Documentation](#) | [Timeline](#) | [Roadmap](#) | **[Browse Source](#)** | [View Tickets](#) | [Open a new ticket](#) | [Search](#) | [Calendar](#) | [Admin](#) | [Gantt chart](#)

[Subversion Location](#)

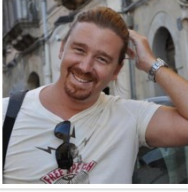
## Repository Index

Name ▲	Size	Rev	Age	Author	Last Change
▼ <a href="#">repo</a>		544	5 months	anonymous	Added 'Makefile' to tarball. Excluded '.svn' files from tarball ...
▶ <a href="#">branches</a>		2	3 years	root	initial import
▶ <a href="#">prod</a>		482	15 months	anonymous	added a new branch for production code
▶ <a href="#">tags</a>		38	3 years	anonymous	tagging for the potch tutorial
▶ <a href="#">trunk</a>		540	6 months	anonymous	
▶ <a href="#">users</a>		544	5 months	anonymous	Added 'Makefile' to tarball. Excluded '.svn' files from tarball ...
▶ <a href="#">software_installation</a>		-	-	-	TracError: GIT version 1.5.4.7 installed not compatible(need >= 1.5.6)





# Recognise individuals... wider trust base



**Bruce Becker**  
brucellino

Meraka Institute, CSIR  
Pretoria, South Africa  
bbecker@csir.co.za  
<http://www.sagrid.ac.za>  
Joined on Aug 08, 2012

1 follower   7 starred   8 following

Contributions   Repositories   Public Activity   [Edit Your Profile](#)

### Popular repositories

- [SAGridIdP](#)  
Code for the Catch-All IdP of South Africa   1 ★
- [SAGridDemo](#)  
A repository of demo applications to sho...   1 ★
- [SAGridApps](#)  
SAGrid Applications   0 ★
- [Quarks-Matter](#)  
Quarks Matter - the story of how and why.   0 ★
- [test](#)  
playground   0 ★

### Repositories contributed to

- [SAGridOps/SoftwareInstallation](#)  
The installation scripts for SAGrid Applic...   1 ★
- [SAGridOps/ServiceDocumenta...](#)  
A repository to contain the documentati...   0 ★
- [SAGridOps/Documents](#)  
A place for public documents   0 ★


### Your Contributions

Summary of Pull Requests, issues opened and commits. [Learn more.](#)

Less  More







# Decentralise the team



## SAGrid Operations team

SAGridOps

South Africa  
sagrid-ops@lists.sagrid.ac.za  
http://www.sagrid.ac.za  
Joined on Jul 05, 2013

6 public repos 0 private repos 1 member



Repositories Members

Edit SAGridOps's Profile

Find a repository... Search

All Public Private Sources Forks Mirrors

**grid-services-deployment** Shell ★ 0 📄 1  
 forked from auth-scc/grid-services-deployment  
 Example ansible playbooks that are used to deploy Grid services  
 Last updated 12 days ago

**Documents** ★ 0 📄 0  
 A place for public documents  
 Last updated 16 days ago

**DevOps** ★ 0 📄 0  
 Top-level directory of devops scripts for SAGrid  
 Last updated 20 days ago

**SoftwareInstallation** Shell ★ 1 📄 0  
 The installation scripts for SAGrid Application deployment.  
 Last updated 2 months ago

**ServiceDocumentation** ★ 0 📄 0  
 A repository to contain the documentation and planning for various services. Includes technical documentation, prototypes, etc.



# Collaborate with code



Repositories Members [Edit SAGridOps's Profile](#)

Find a repository...   [All](#) [Public](#) [Private](#) [Sources](#) [Forks](#) [Mirrors](#)

**grid-services-deployment** Shell ★ 0 📄 1  
forked from [auth-scc/grid-services-deployment](#)  
Example ansible playbooks that are used to deploy Grid services  
Last updated 12 days ago

**SAGrid Operations team**  
SAGridOps

South Africa  
✉ [sagrid-ops@lists.sagrid.ac.za](mailto:sagrid-ops@lists.sagrid.ac.za)  
🌐 <http://www.sagrid.ac.za>  
🕒 Joined on Jul 05, 2013

6 public repos 0 private repos 1 member

**Documents** ★ 0 📄 0  
A place for public documents  
Last updated 16 days ago

**DevOps** ★ 0 📄 0  
Top-level directory of devops scripts for SAGrid  
Last updated 20 days ago

**SoftwareInstallation** Shell ★ 1 📄 0  
The installation scripts for SAGrid Application deployment.  
Last updated 2 months ago

**ServiceDocumentation** ★ 0 📄 0  
A repository to contain the documentation and planning for various services. Includes technical documentation, prototypes, etc.



## Let the robots do the work

- Define what we want to deploy – let the experts take care of how to deploy
- **DevOps** paradigm – same review/tag/release mechanisms on operations code as we have for scientific applications
  - Teach a marketable skill
  - Allow specialisation
  - Enable remote management of complex services
  - Ensure that published methodology is adopted methodology





# Quality Control and feedback



- Ensure that requested applications are included in the repo
- Provide testing and QA infrastructure
- Self-serve to users

***Talk to Fanie Riekert***  
**fanie.riekert@gmail.com**

# Well done, you made it to the end of the talk



- "SAGrid" is now a misnomer – have you got a suggestion ?
  - NGI\_ZA + lots of other services in a federation
  - A lot of assumptions about *what* and *how* need to be updated
- Underlying infrastructure and core services need to be reliable – measured and improved
- Peer MoU with EGI.eu
  - greatly increased reliability/availability and scope of services
  - Direct collaboration with other NGI/NRENS
- Eat your own dog food – if it tastes bad, don't be afraid to change it
- Remove barriers to access, usage and development
- Predicting, reporting, managing, implementing infrastructure services



# Thank you !

- It's a community... there's no "us" and "you"
  - SAGrid on Facebook and Twitter
    - <https://www.facebook.com/SAGrid>
    - <https://www.twitter.com/TheSAGrid>
- There's place for everyone and there doesn't necessarily have to be only 1 leader

