Hunting The Higgs

Using the Worldwide LHC Computing Grid (WLCG)

Abstract

- This talk describes the long process of hardening the WLCG through a series of Service Challenges and production readiness tests leading to reliable petascale computing services that allowed data to be turned into discoveries in record time.

- It covers not only the preparation period for Run1 of the LHC, the data-collection period that led to the announcement of the discovery of a Higgs-like particle, but also the preparations for Run2 and beyond – higher data rates, more flexible network architectures and the challenges of tomorrow's processors.

- It also describes collaborative work performed in the context of the EGI-InSPIRE Heavy User Communities work package in developing and sharing common solutions across multiple communities.

Caveat

- It is not my intention to give a purely historical talk but rather try to understand how future challenges – such as the SKA – can be met
- After all, the only thing that we learn from history is that we never learn from history

LHC Computing – A Grand Challenge

- In 1978, a former DG of CERN proposed that the LEP tunnel – should it be built – be large enough for a future hadron collider (on top)
- It was in September 1992 that I realized I was working essentially full-time on LHC Computing
- Daily WLCG Operations meetings started in 2008: first collisions of accelerated pp in 2010;
- "Higgs day" – 4th July 2012
- But we are still just at the beginning….

1990s: Revolution(s) in HEP Computing

- In parallel with the exploitation of LEP, more or less every “constant” in terms of HEP Computing was over-turned
- Arrival of networking (multiple protocols)
- End of the main-frame: minis and micros were rapidly superseded by PCs – and Linux
- Fortran – plus extensions – were replaced by C++
- Numerous R&D projects: ODBMS, re-implementation of key HEP s/w packages…
- MONARC: “the set of viable computing models”
- “622 Mb/s to a few centres with tape as backup”

2000s: “Enter the Grid”

- With x86, Linux and C++ as the new norm, HEP became infected with grid-itus in the early 2000s
- Assisted by a series of EU projects (EDG, EGEE I/II/II, EGI) with other efforts in the US and elsewhere, production grid computing appeared
- But it was not without significant effort (blood, sweat, tears)
- Now that “it works”, many have already forgotten what it took to get there

2004 – Service Challenges Begin

- Originally conceived as a series of 4 challenges, later supplemented by two “production readiness” tests
- SC1 & 2 did not involve users (and failed to meet targets)
- SC3 failed to meet targets but a re-run 6 months later was successful
- Step-by-step pains-taking debugging
- Even SC4 did not include the “analysis use case” – only successfully demonstrated 2 years after target startup-date of 2007

Reminder – one of the conclusions from the plenary talk at CHEP' 04 by Fabiola Gianotti

My 2 main worries today (as an LHC physicist and end-user):

- End-users not yet exposed to massive use/navigation of database and of GRID
  → what will happen when $O(10^3)$ physicists will simultaneously access these systems ?
- Software and Computing Model developed for steady-state LHC operation (≥ 2009 ?)
  But : at the beginning they will be confronted with most atypical (and stressful) situations,
  for which a lot of flexibility will be needed:
  -- staged, non-perfect, non-calibrated, non-aligned detectors with all sorts of problems
  -- cosmic and beam-halo muons used to calibrate detectors during machine commissioning
  -- machine backgrounds ; higher-than-expected trigger rates
  -- fast/frequent reprocessing of part of data (e.g. special calibration streams)
  -- $O(10^3)$ physicists in panic-mode using and modifying the Software and accessing the database, GRID ...

⇒ it is time for the Software/Computing to address the early phase of LHC operation, not to hinder the fast delivery of physics results (and a possible early discovery ...)

The LCG Service Challenges:

Rolling out the LCG Service

Jamie Shiers, CERN-IT-GD-SC

http://agenda.cern.ch/fullAgenda.php?ida=a053365

June 2005

LCG Service Challenges - Overview

- LHC will enter production (physics) in summer 2007
  - Will generate an enormous volume of data
  - Will require huge amount of processing power
- LCG 'solution' is a world-wide Grid
  - Many components understood, deployed, tested..
- But…
  - Unprecedented scale
  - **Humungous challenge of getting large numbers of institutes and individuals, all with existing, sometimes conflicting commitments, to work together**
- LCG must be ready at full production capacity, functionality and reliability in **little more than 1 year** from now
  - Issues include h/w acquisition, personnel hiring and training, vendor rollout schedules etc.
- **Should not limit ability of physicist to exploit performance of detectors nor LHC's physics potential**
  - Whilst being stable, reliable and easy to use

LCG Service Hierarchy

Tier-0 – the accelerator centre

- Data acquisition & initial processing
- Long-term data curation
- ➢ **Data Distribution to Tier-1 centres**

Tier-1 – "online" to data acquisition process → high availability

- Managed Mass Storage – → grid-enabled data service
- ➢ **All re-processing passes**
- Data-heavy analysis
- National, regional support

Tier-2 – ~100 centres in ~40 countries

- Simulation
- End-user analysis – batch and interactive
- ➢ **Services, including Data Archive and Delivery, from Tier-1s**

les.robertson@cern.ch

Tier 0 at CERN: Acquisition, First pass reconstruction, Storage & Distribution

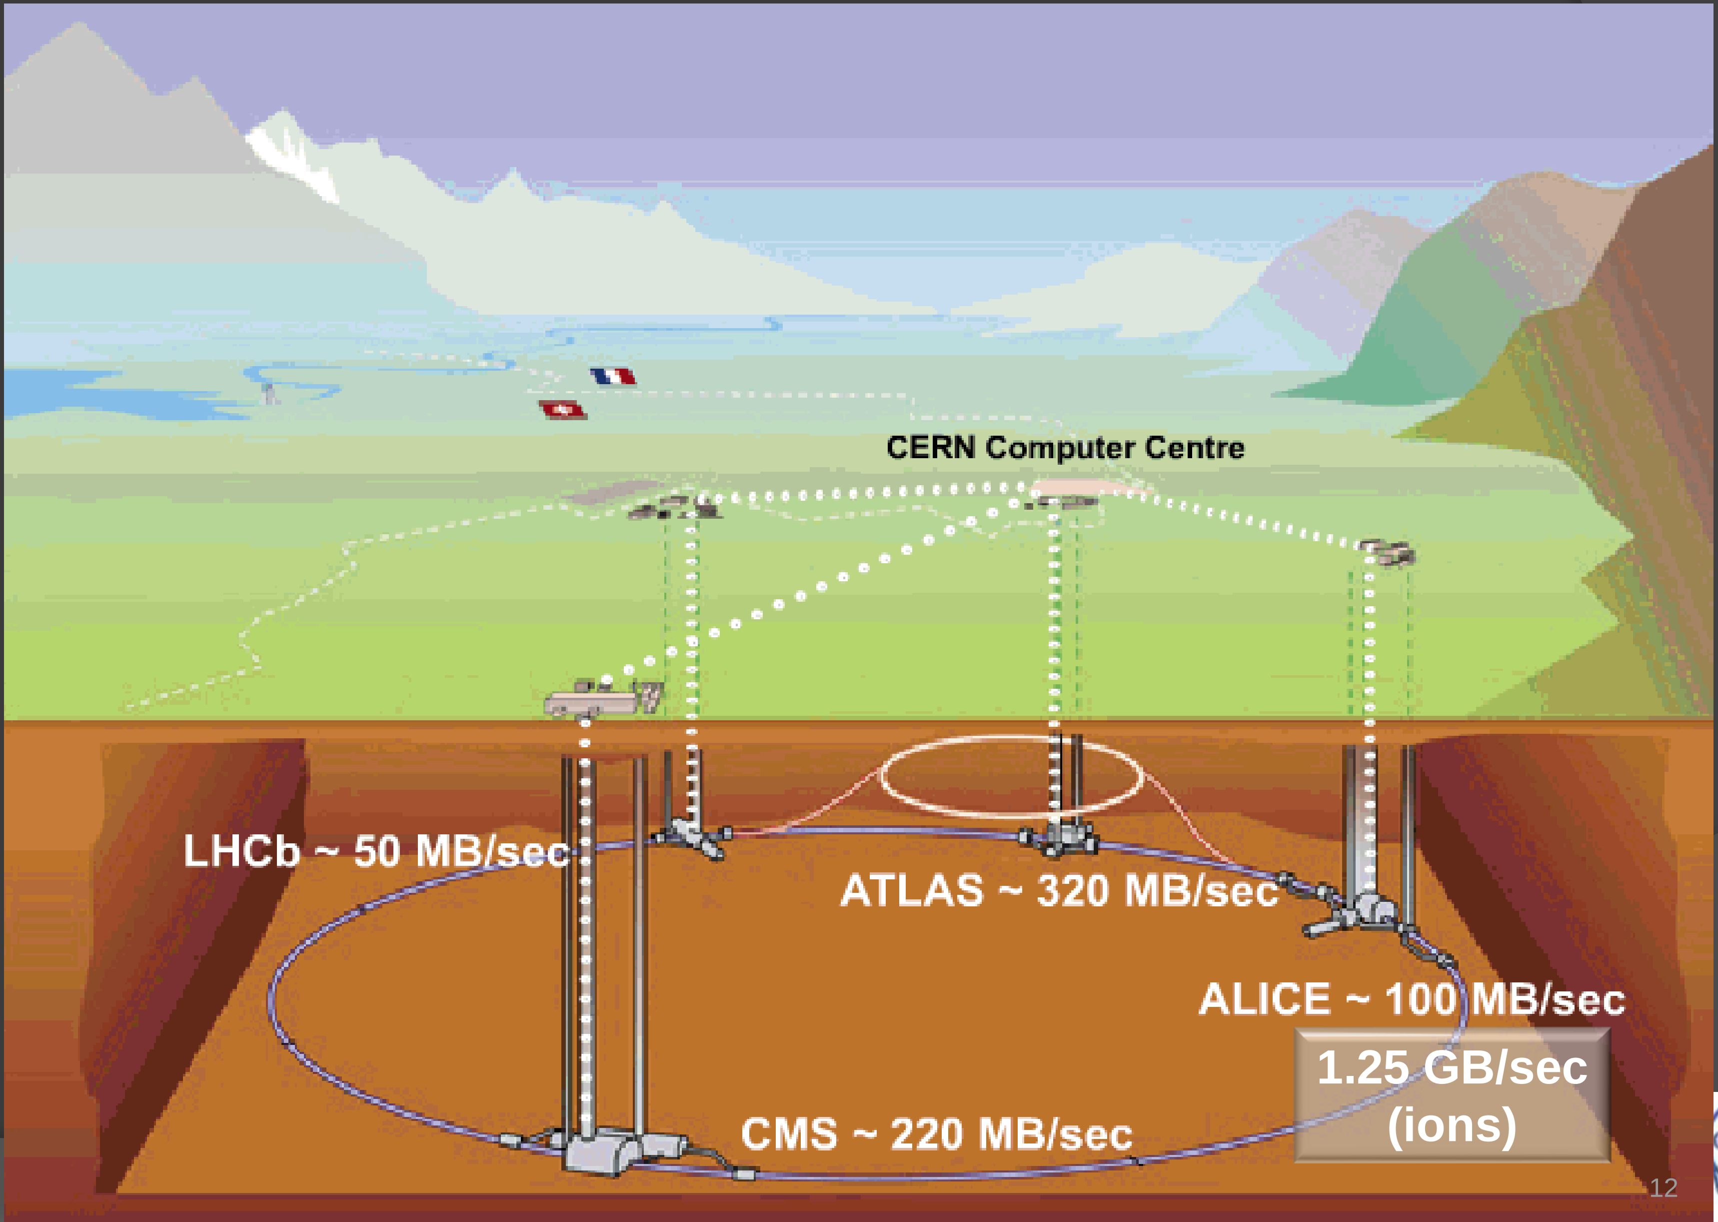

Tier 0 – Tier 1 – Tier 2

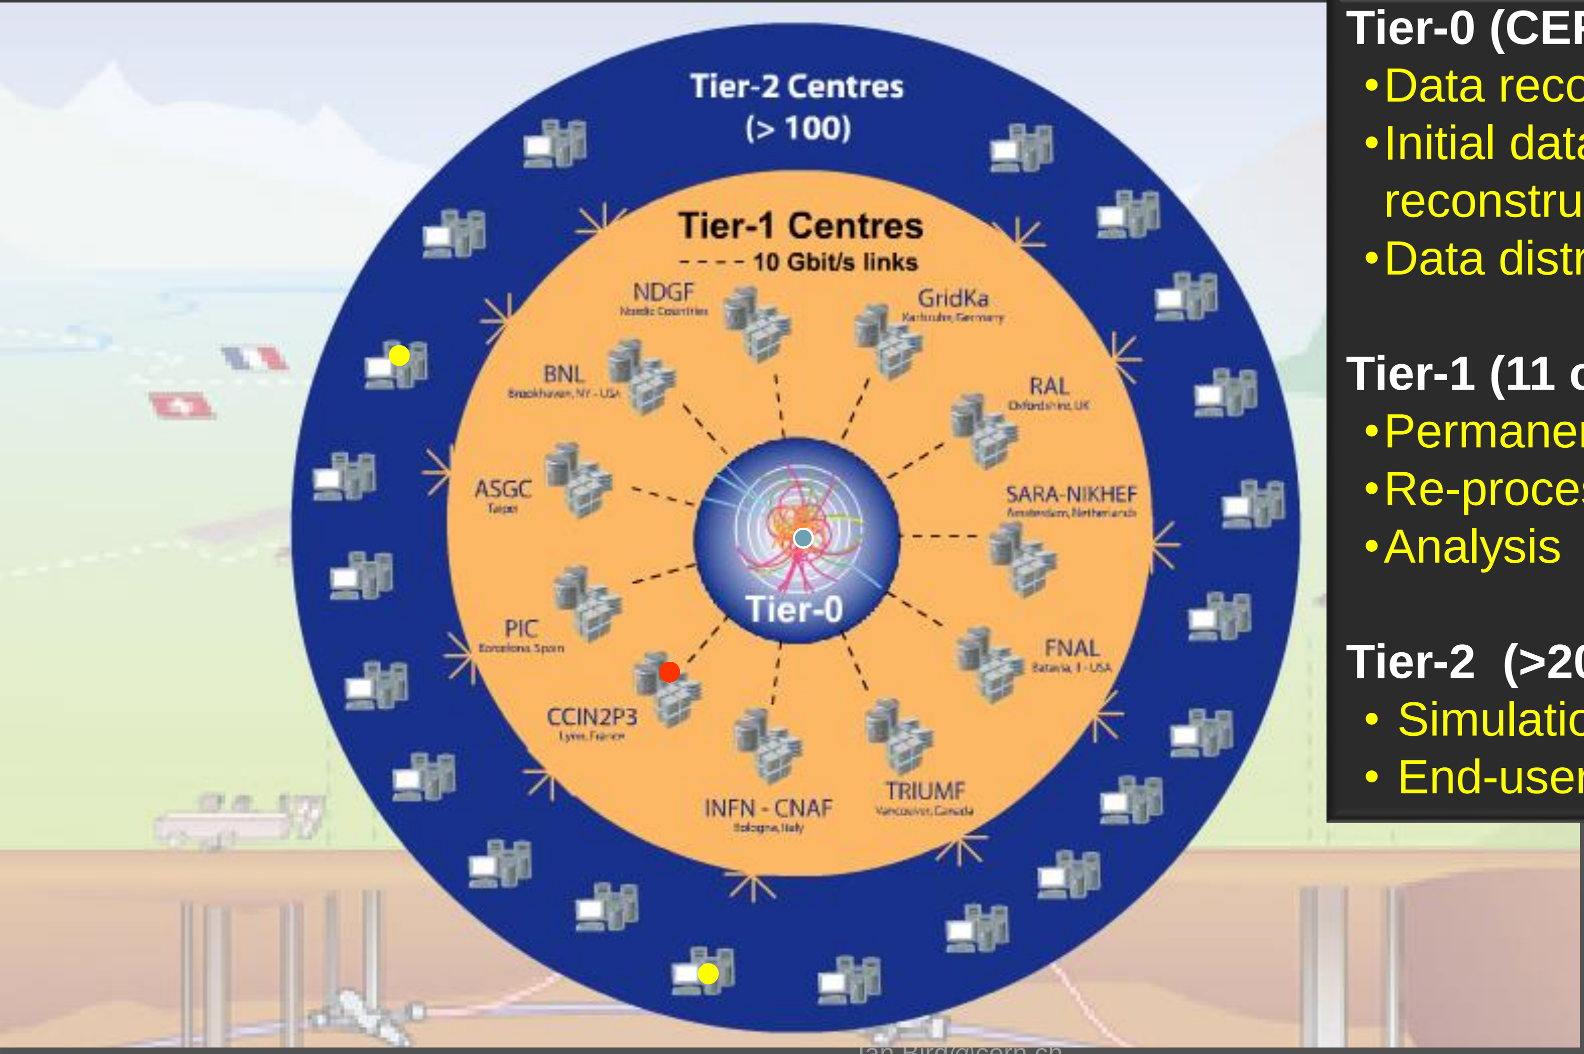

Tier-0 (CERN):
- Data recording
- Initial data reconstruction
- Data distribution

Tier-1 (11 centres):
- Permanent storage
- Re-processing
- Analysis

Tier-2 (>200 centres):
- Simulation
- End-user analysis


BNL
ASGC/Taipei
CCIN2P3/Lyon
TRIUMF/BC
NIKHEF/
SARA
FNAL
RAL
PIC
NDGF
CNAF
TIER2s
CERN
FZK

LCG Deployment Schedule

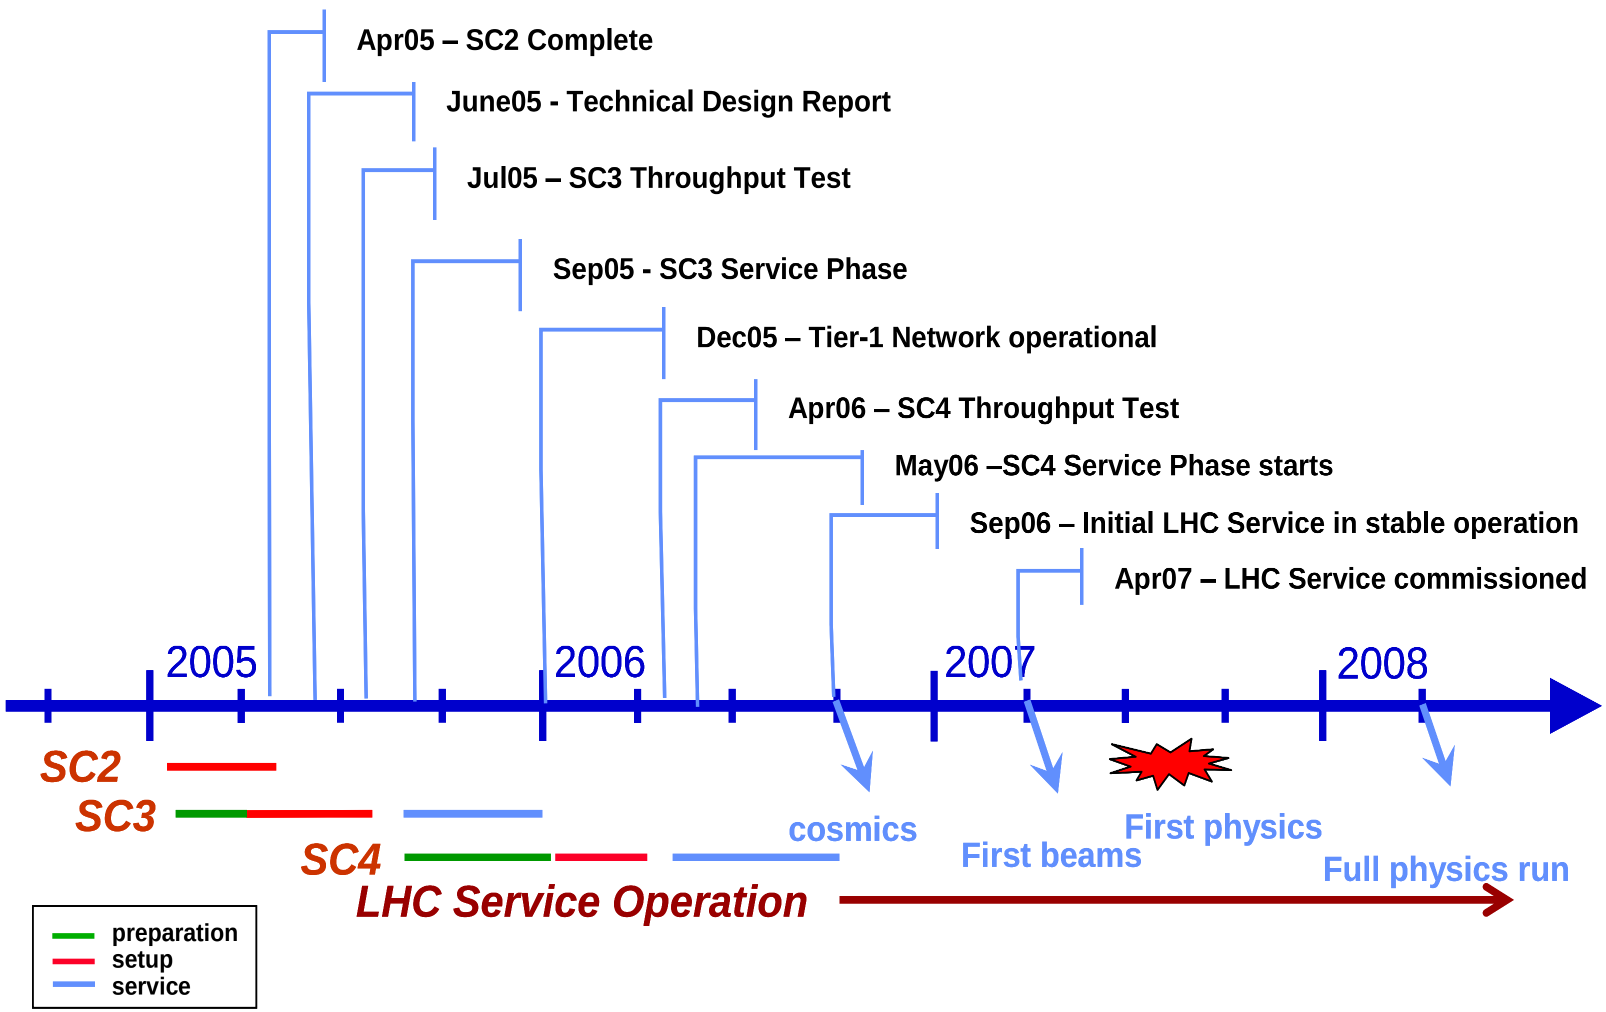

Metrics for Success

- (EU) projects typically have a bunch of metrics
- EGI-InSPIRE had, for example:
  - # GGUS tickets; MTR etc.
- But no clear understanding of what represented “good” or “progress”
- We developed some very simple metrics, still in use today:
  - Examined at “weekly” (now monthly or less) MB meetings;
  - Examined via quarterly reports at CB/OB level
- Plus a “high-level” metric…

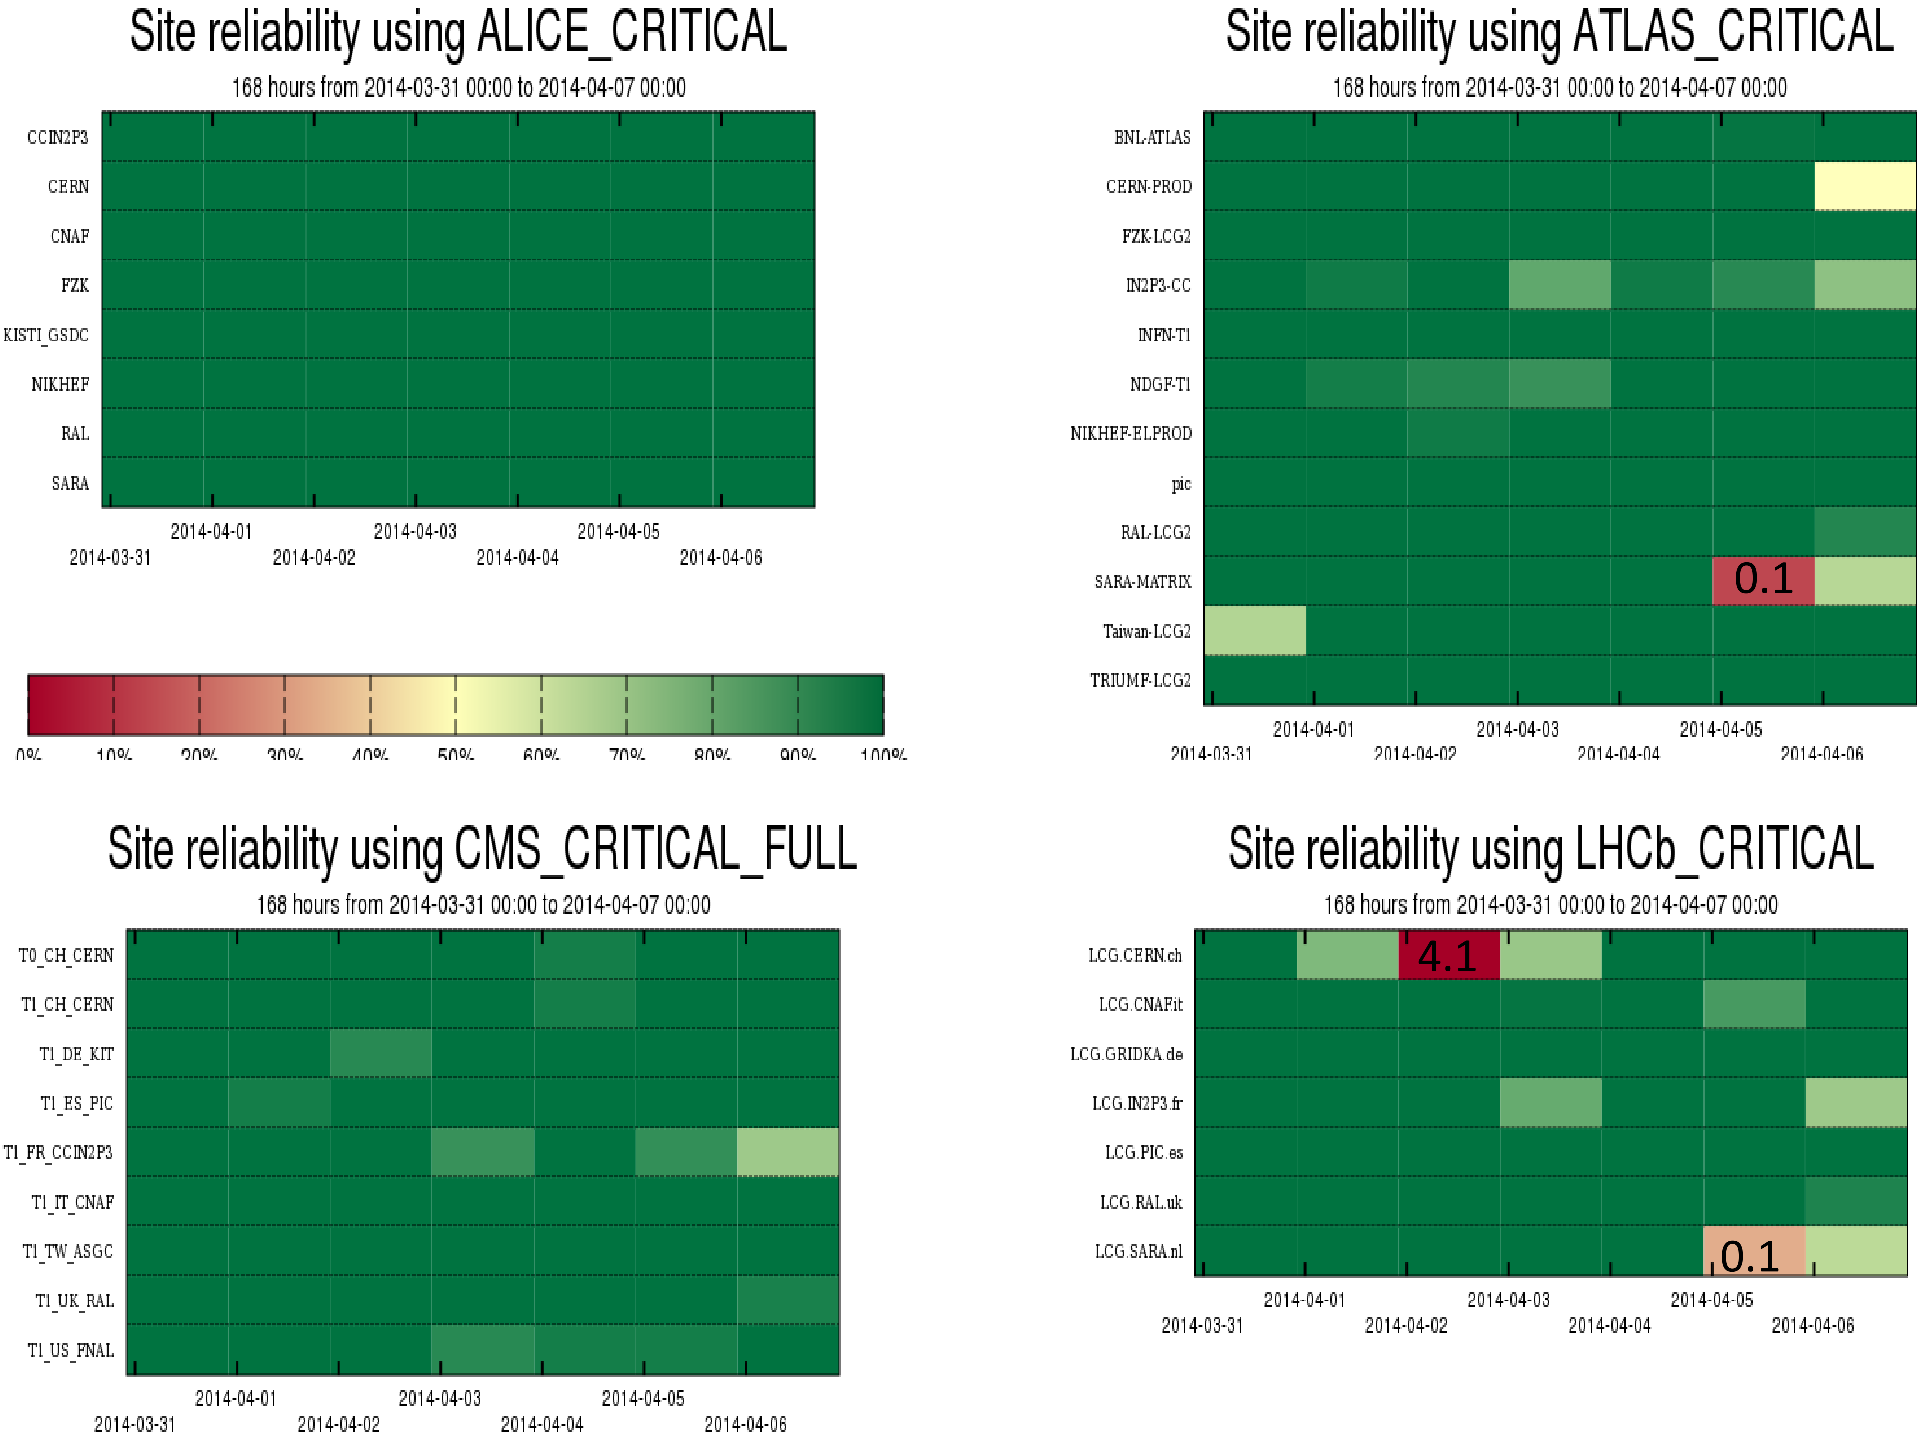

Site reliability using ALICE_CRITICAL
168 hours from 2014-03-31 00:00 to 2014-04-07 00:00
CCIN2P3
CERN
CNAF
FZK
KISTI_GSDC
NIKHEF
RAL
SARA
2014-04-01
2014-04-03
2014-04-05
2014-03-31
2014-04-02
2014-04-04
2014-04-06
0%
10%
20%
30%
40%
50%
60%
70%
80%
90%
100%
Site reliability using ATLAS_CRITICAL
168 hours from 2014-03-31 00:00 to 2014-04-07 00:00
BNL-ATLAS
CERN-PROD
FZK-LCG2
IN2P3-CC
INFN-T1
NDGF-T1
NIKHEF-ELPROD
pic
RAL-LCG2
SARA-MATRIX
Taiwan-LCG2
TRIUMF-LCG2
0.1
2014-04-01
2014-04-03
2014-04-05
2014-03-31
2014-04-02
2014-04-04
2014-04-06
Site reliability using CMS_CRITICAL_FULL
168 hours from 2014-03-31 00:00 to 2014-04-07 00:00
T0_CH_CERN
T1_CH_CERN
T1_DE_KIT
T1_ES_PIC
T1_FR_CCIN2P3
T1_IT_CNAF
T1_TW_ASGC
T1_UK_RAL
T1_US_FNAL
2014-04-01
2014-04-03
2014-04-05
2014-03-31
2014-04-02
2014-04-04
2014-04-06
Site reliability using LHCb_CRITICAL
168 hours from 2014-03-31 00:00 to 2014-04-07 00:00
LCG.CERN.ch
LCG.CNAF.it
LCG.GRIDKA.de
LCG.IN2P3.fr
LCG.PIC.es
LCG.RAL.uk
LCG.SARA.nl
4.1
0.1
2014-04-01
2014-04-03
2014-04-05
2014-03-31
2014-04-02
2014-04-04
2014-04-06

WLCG Service Incidents

- Aka “post-mortems”

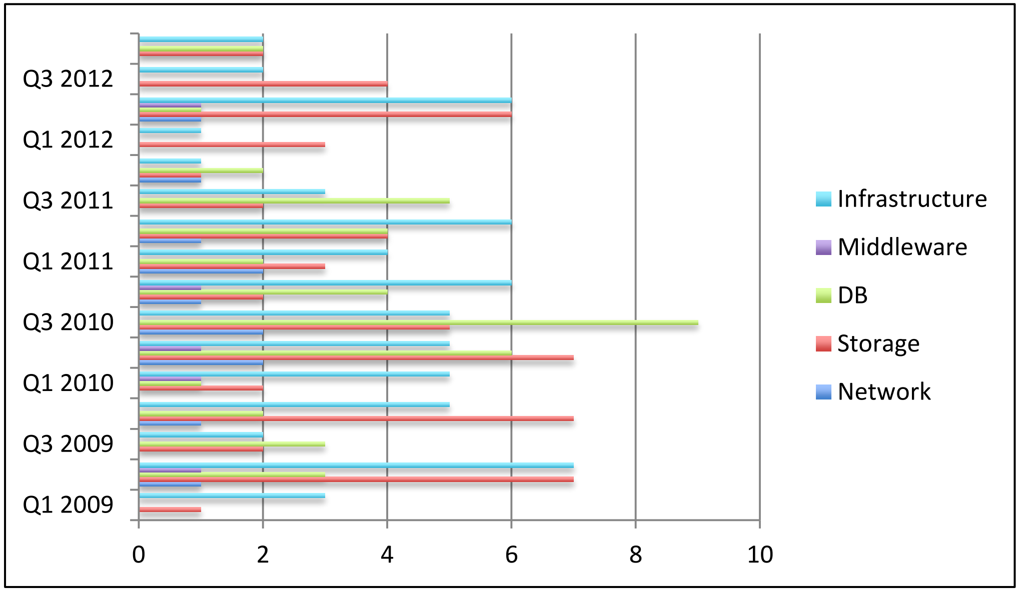

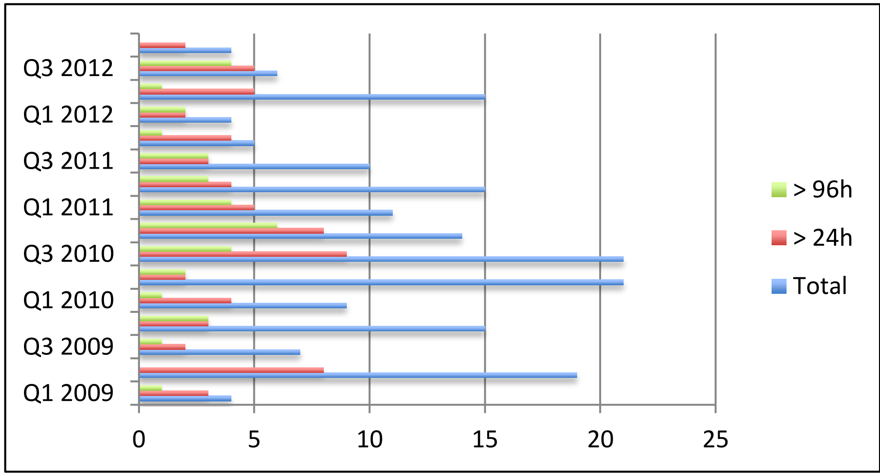

WLCG Service Incidents

These are significant service incidents wrt targets defined in the WLCG MoU. Basically mean major disruption to data taking, distribution, processing or analysis. A Service Incident Report is required.

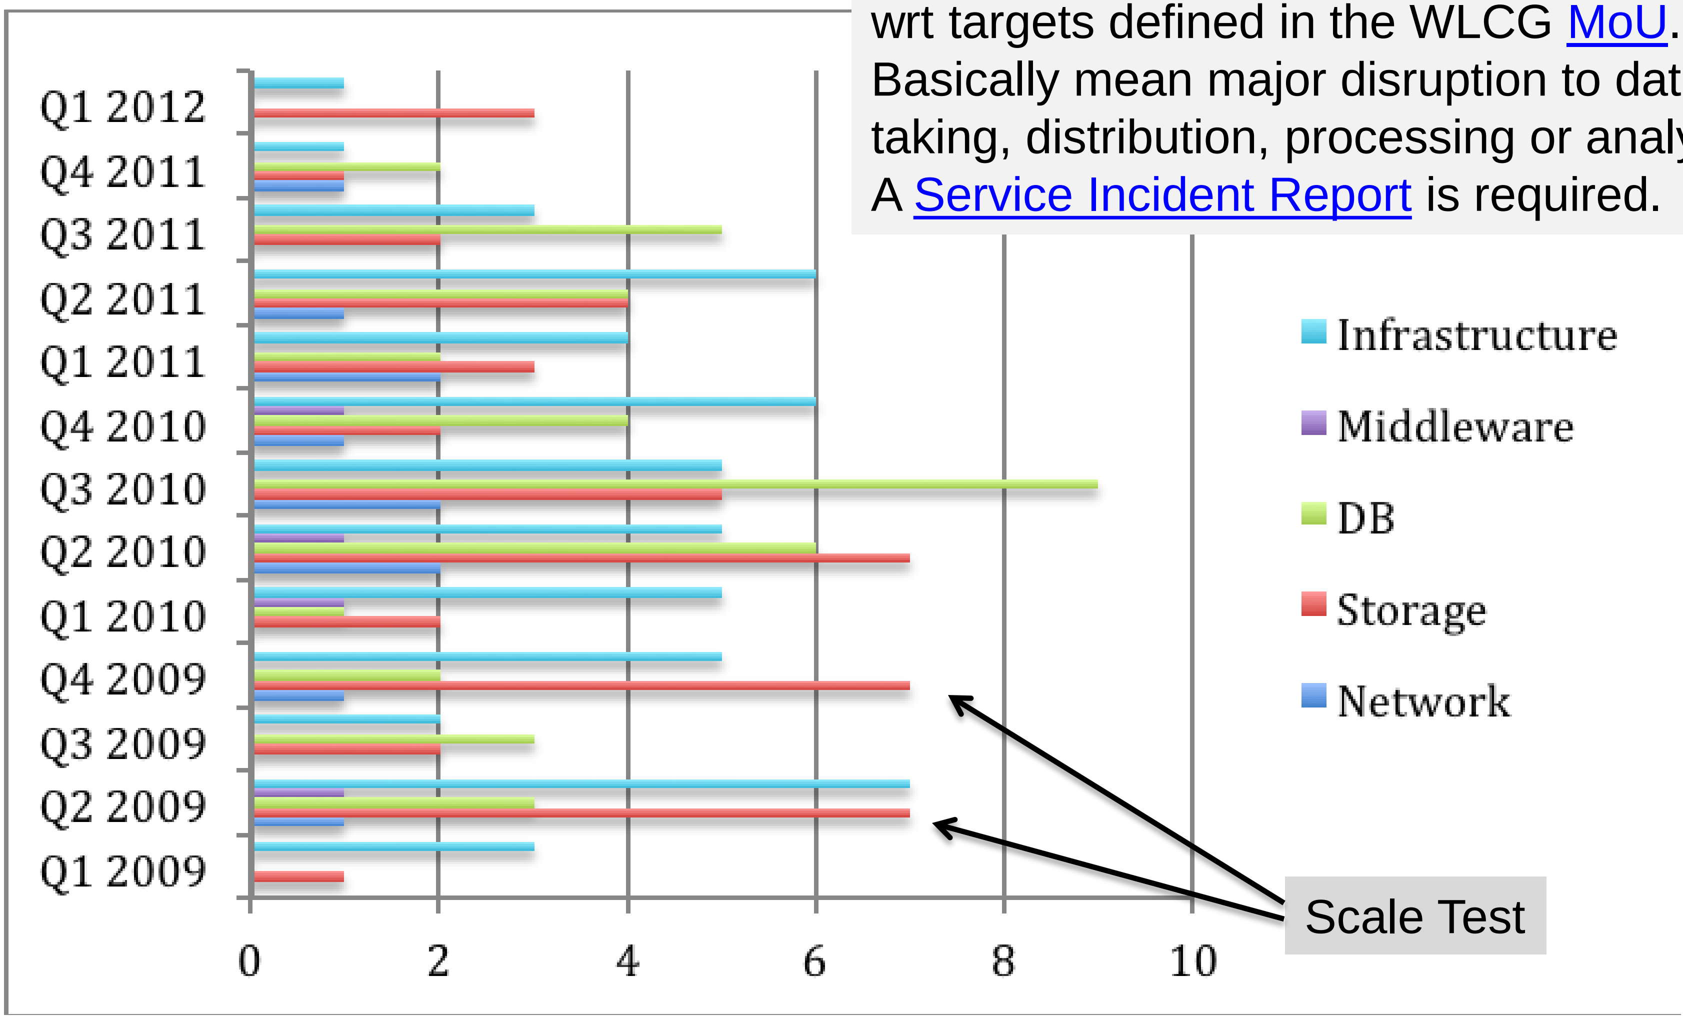

WLCG Service Incidents

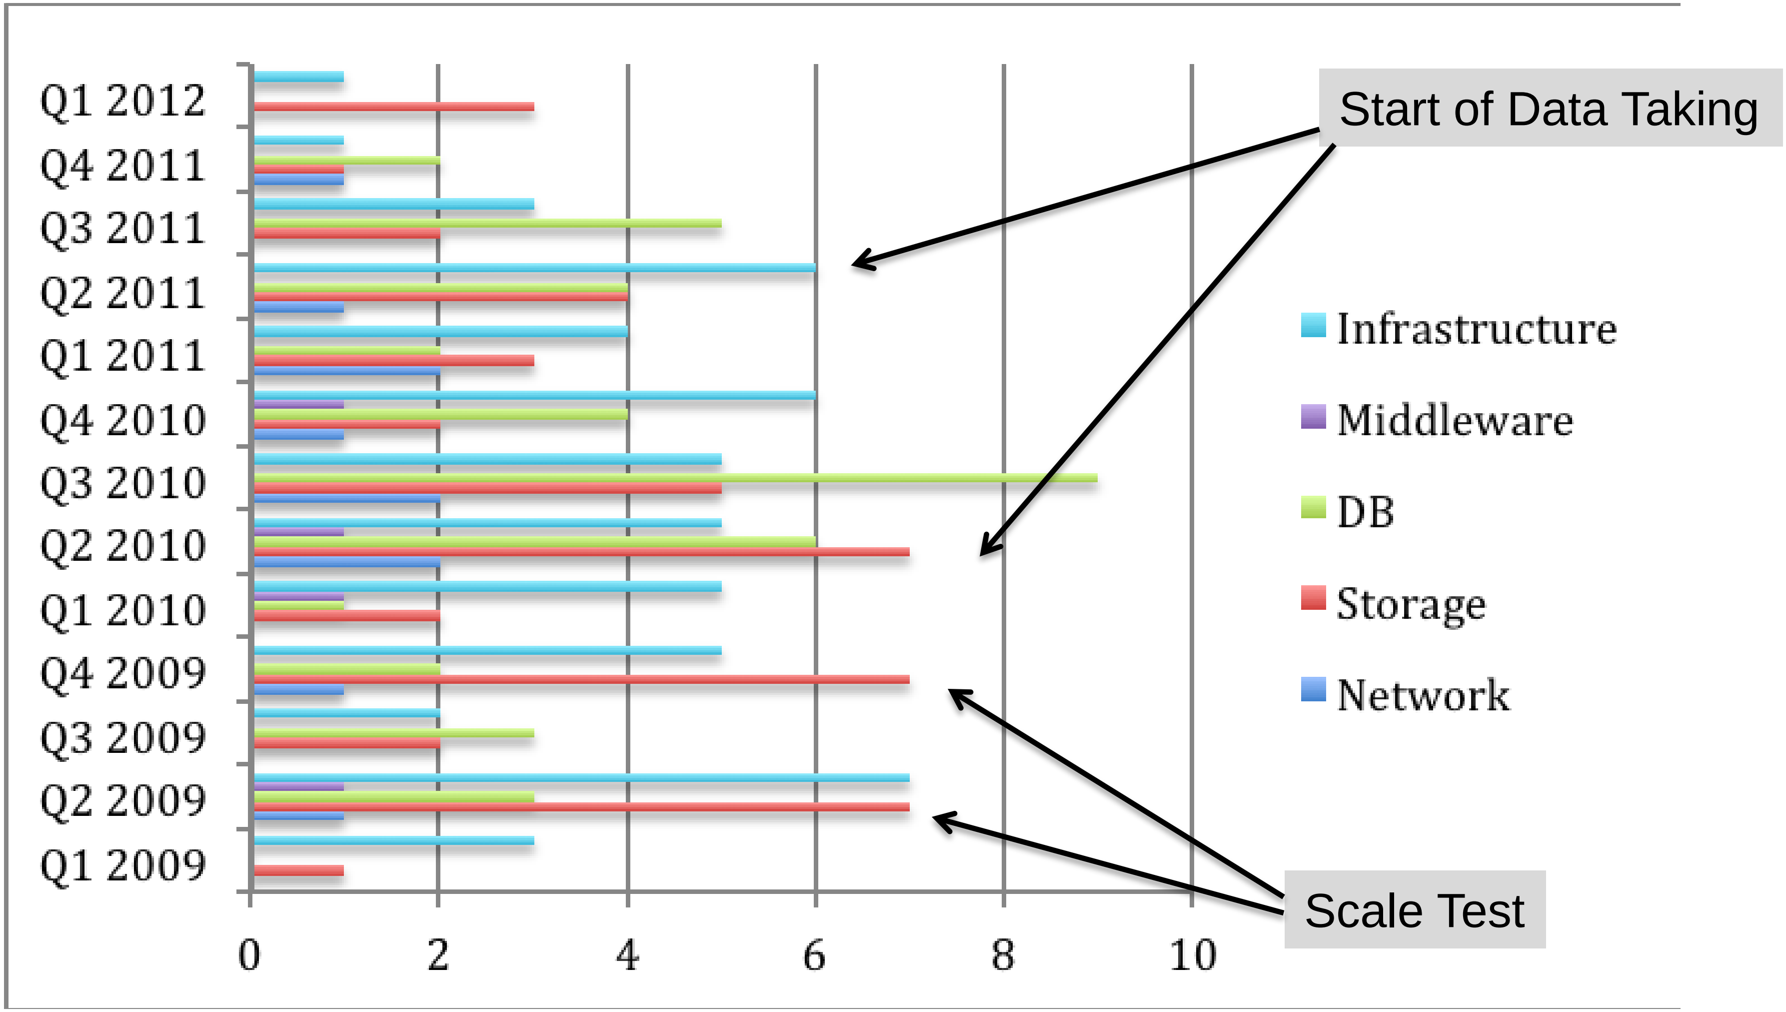

WLCG Service Incidents

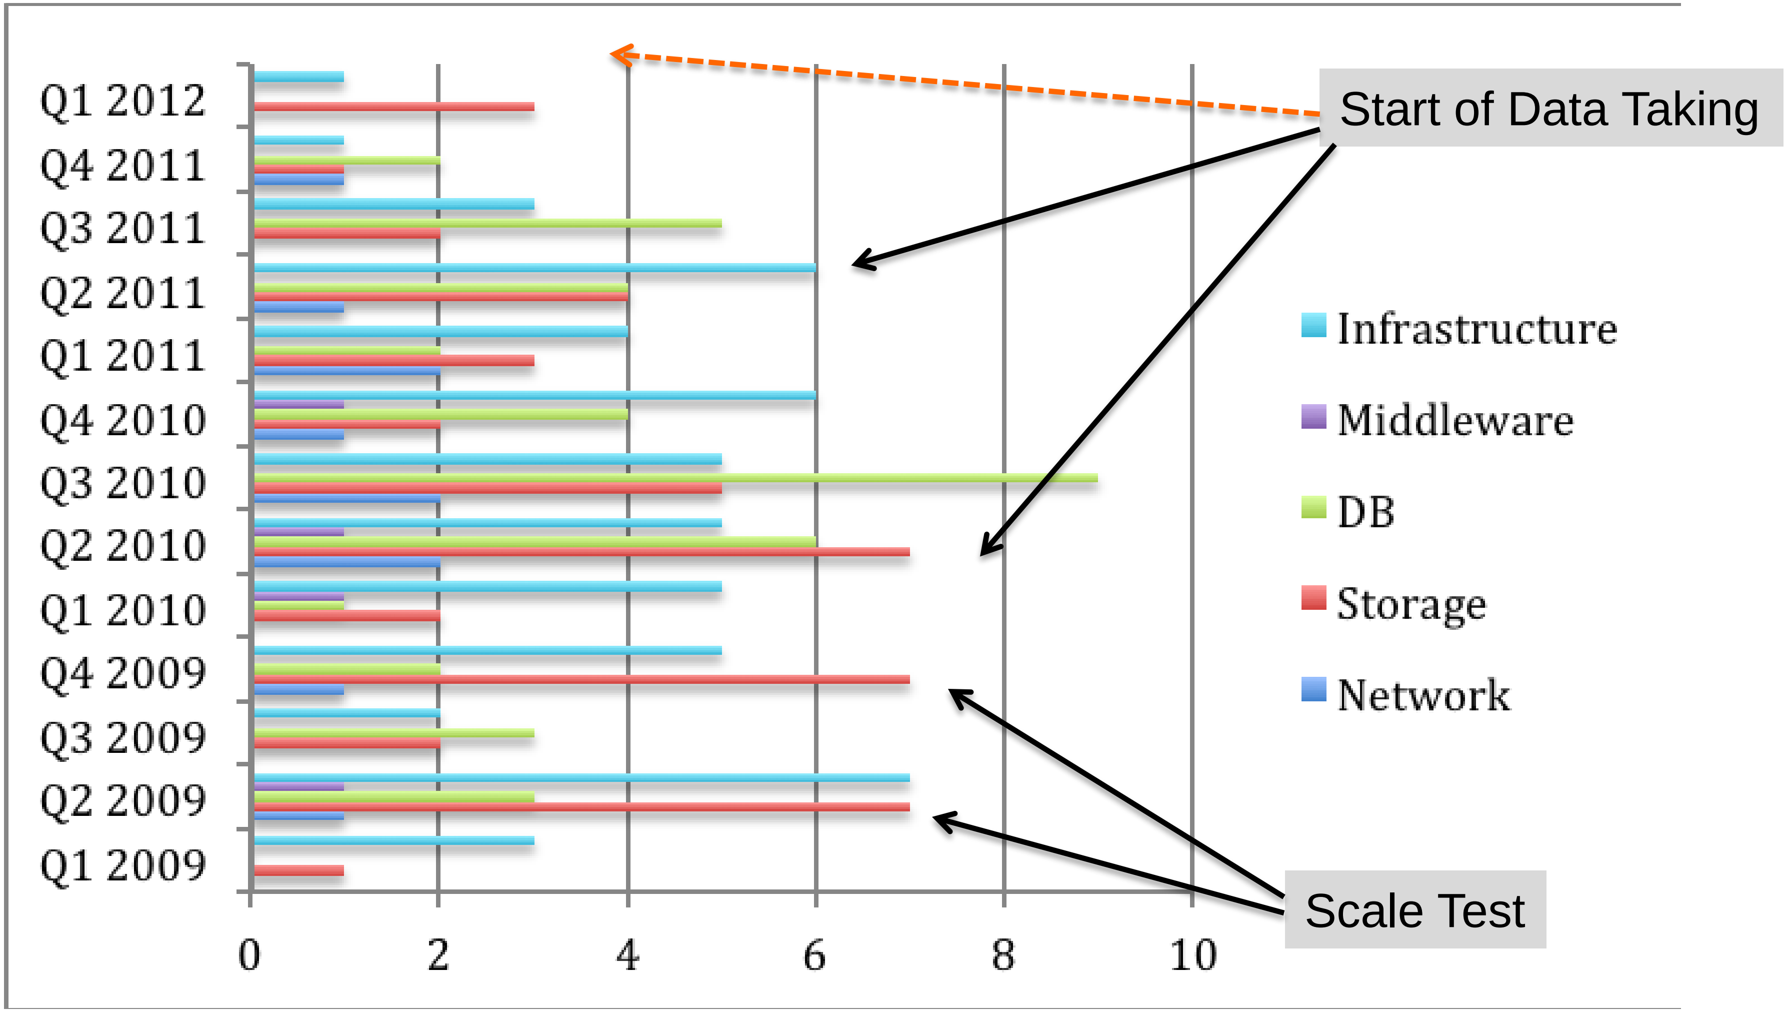

Resolution of Incidents

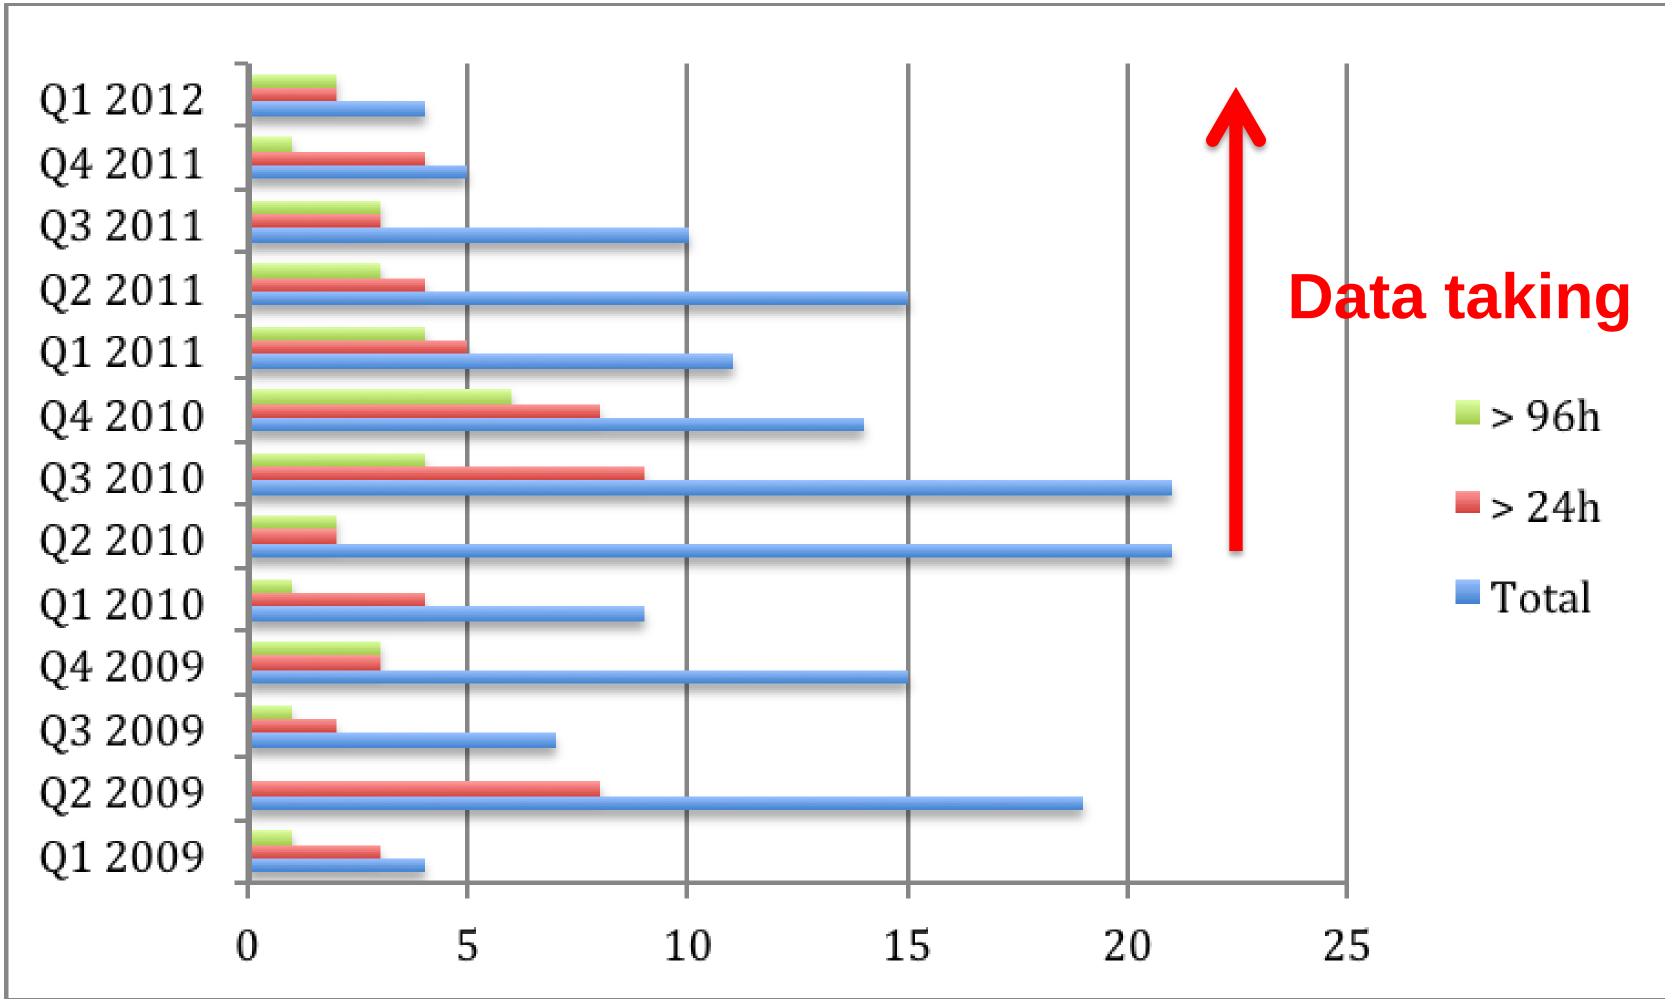

Resolution of Incidents

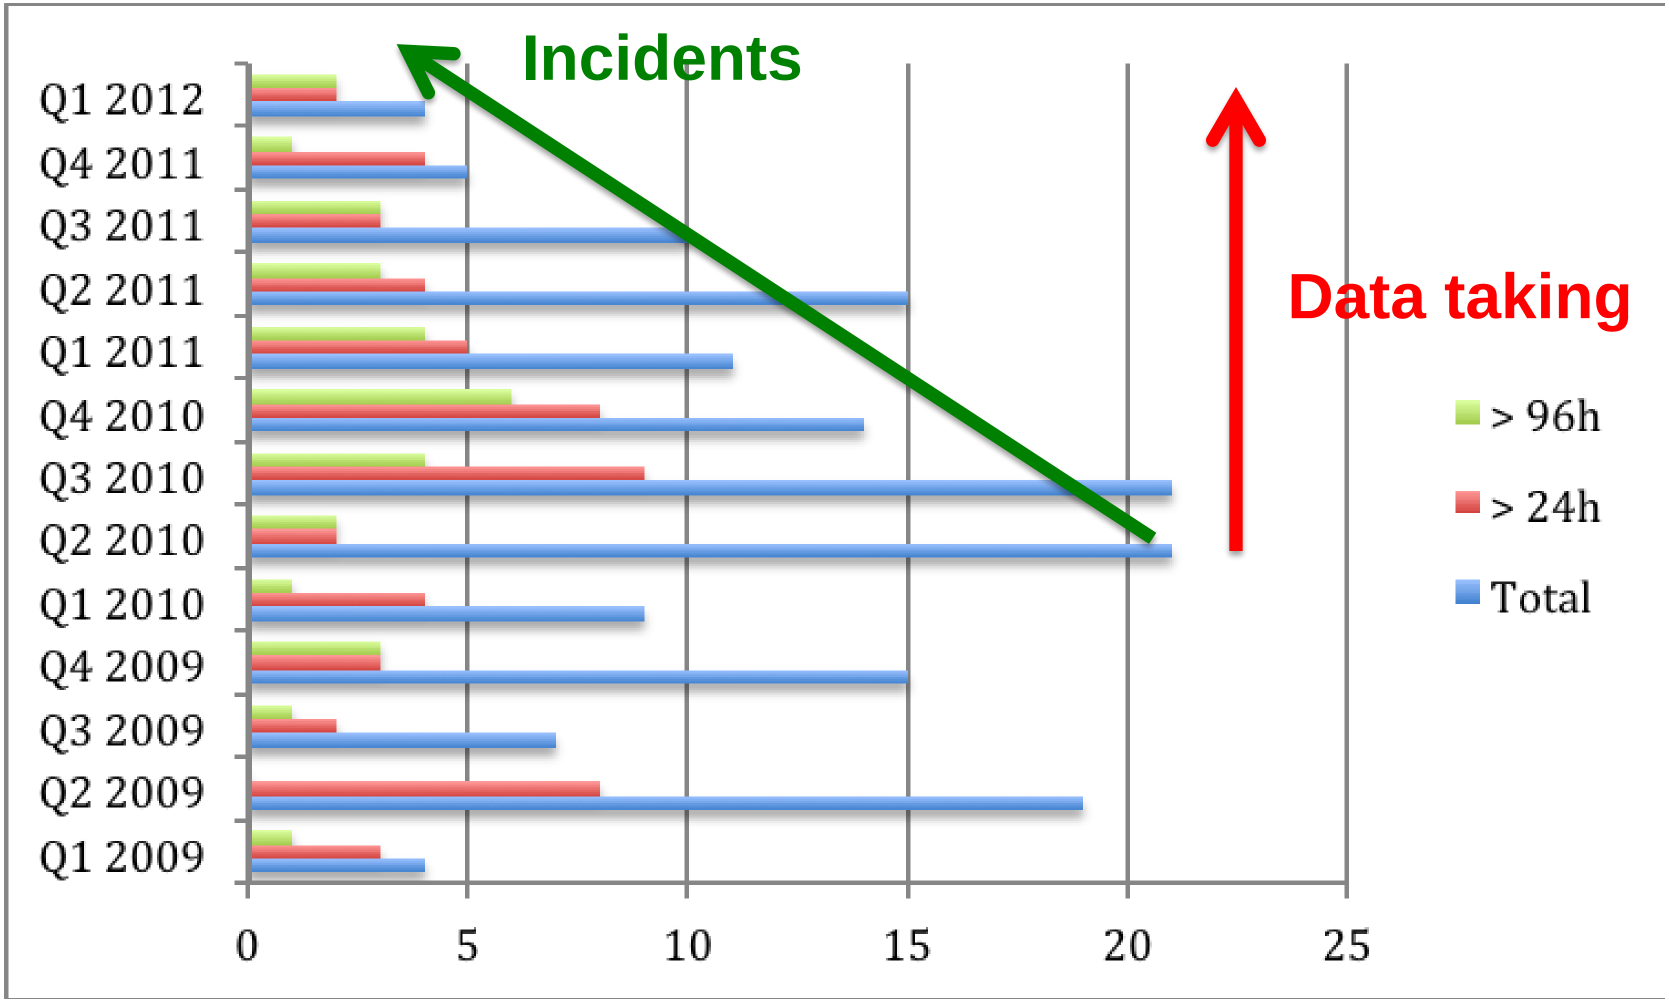

Resolution of Incidents

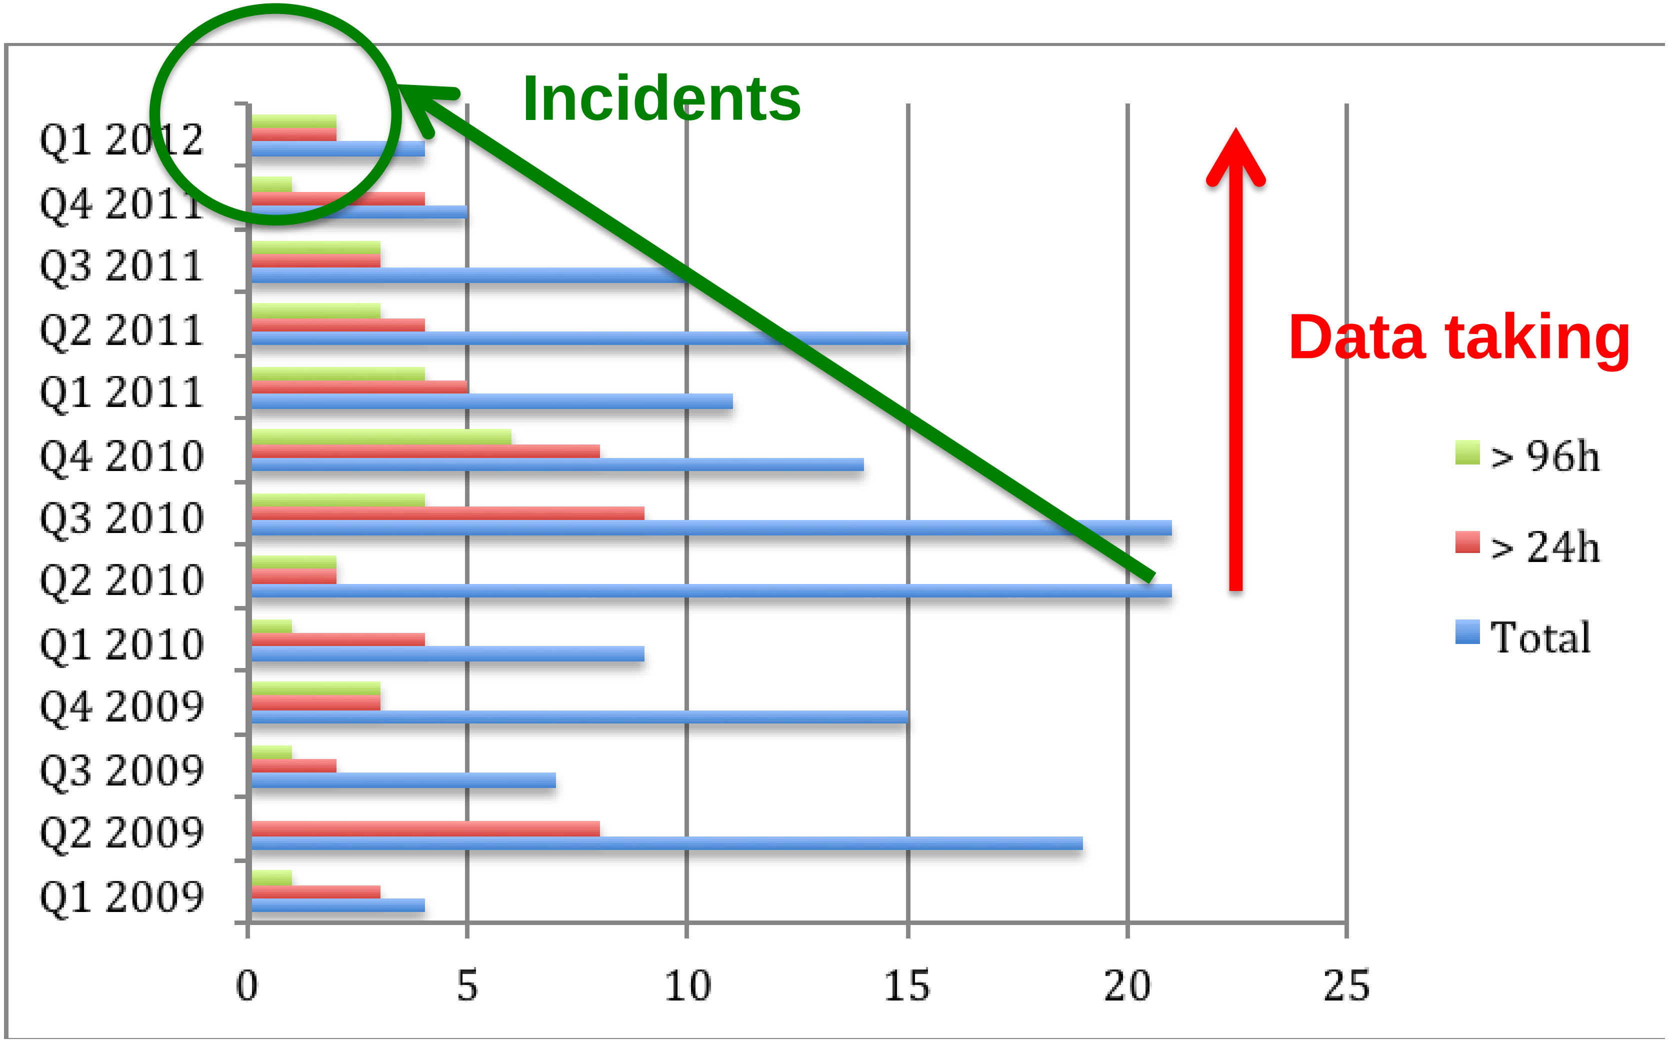

Resolution of Incidents

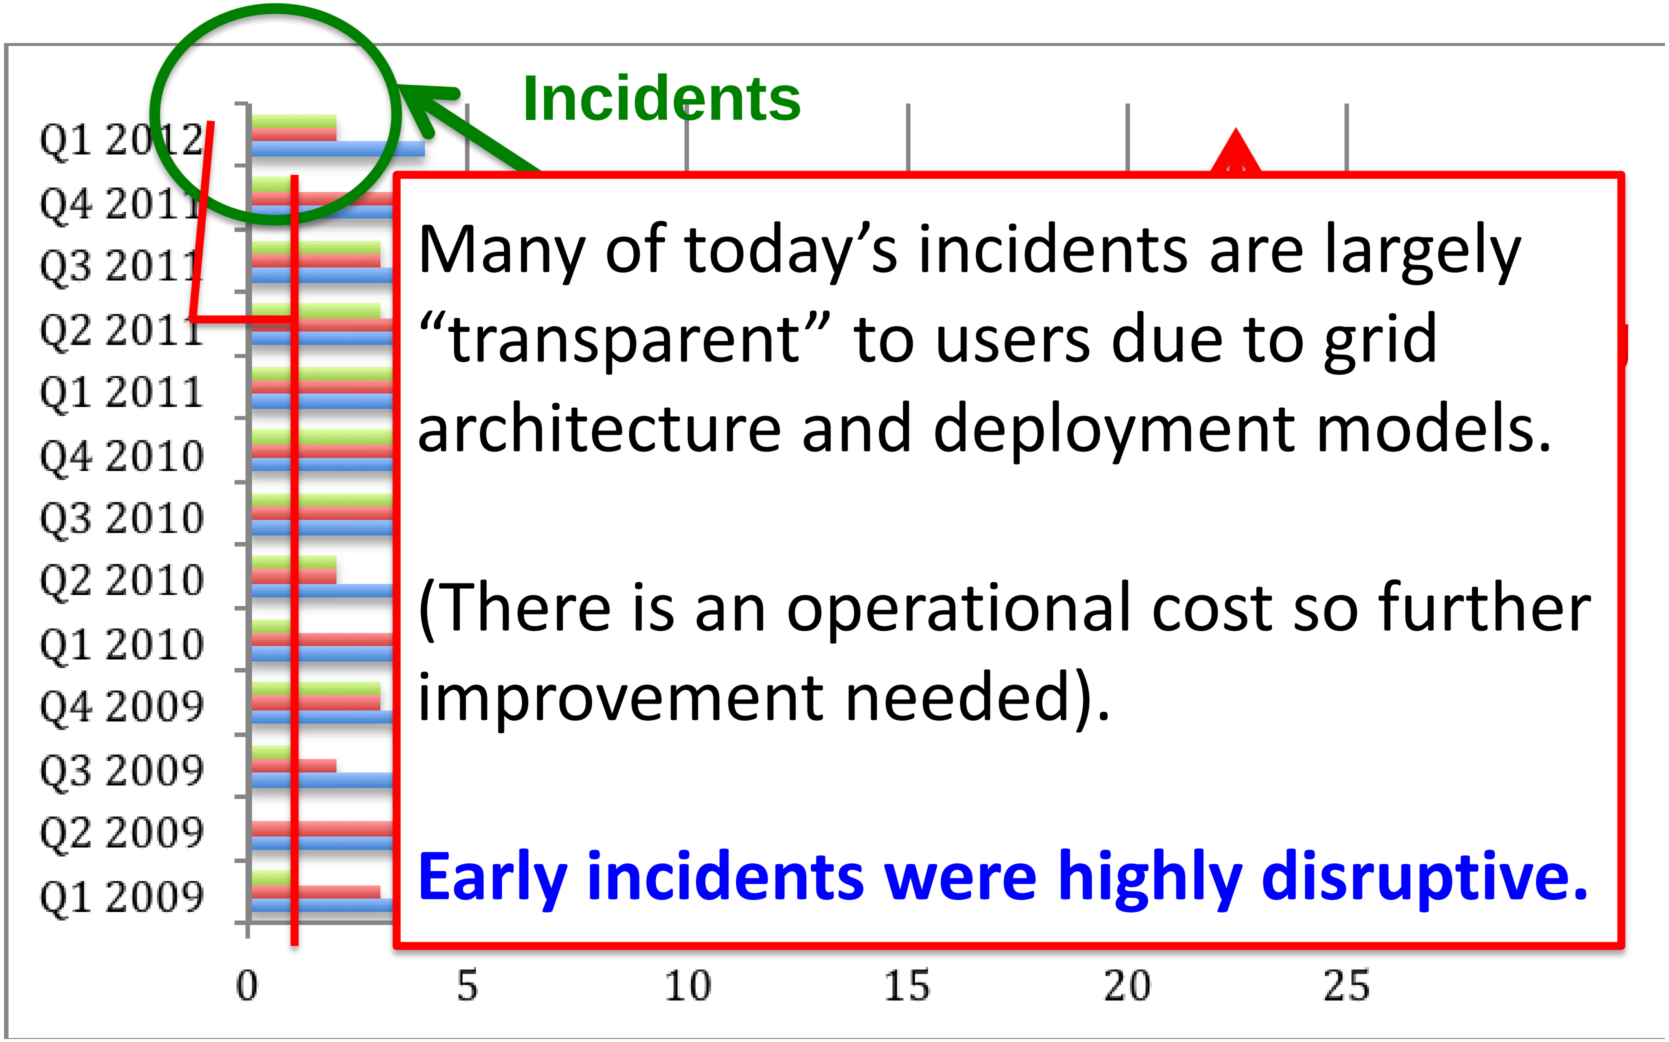

Many of today's incidents are largely "transparent" to users due to grid architecture and deployment models.

(There is an operational cost so further improvement needed).

Early incidents were highly disruptive.

Resolution of Incidents

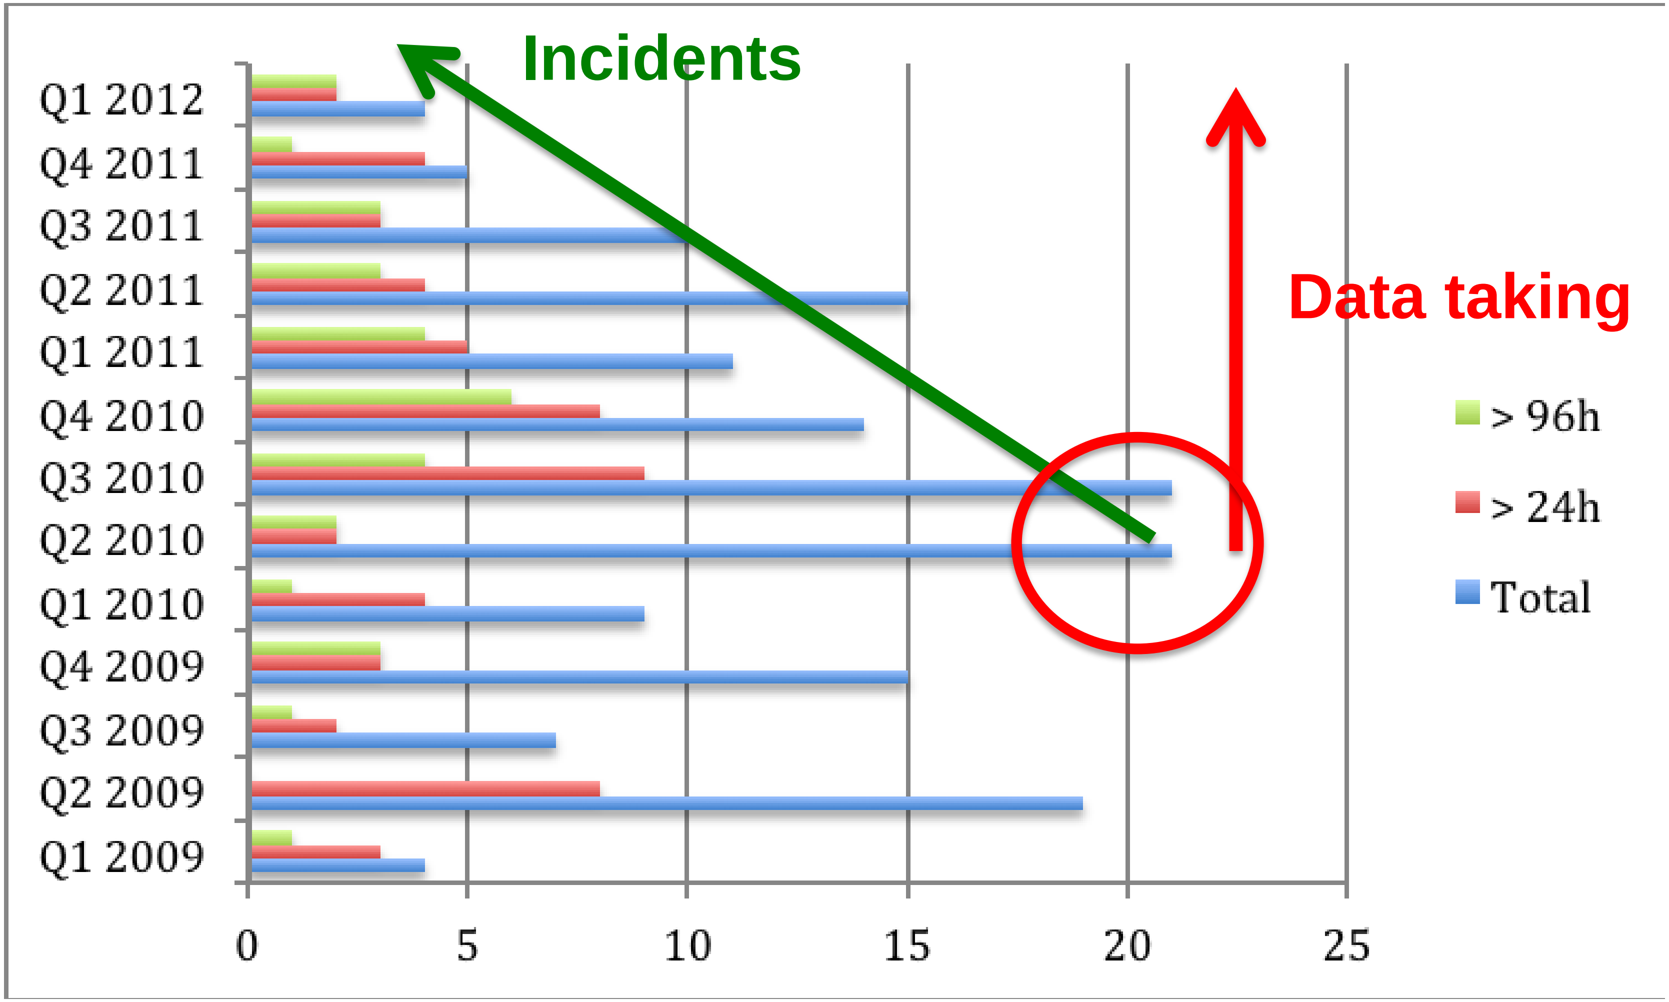

SA3 Overview

9 Countries
11 Beneficiaries
672 PMs
56 FTEs

WP	Task	Beneficiary	Total PMs
WP6-G	TSA3.1	CERN	18
WP6-G	TSA3.2	ARNES	3
WP6-G	TSA3.2	CERN	120
WP6-G	TSA3.2	CNRS	30
WP6-G	TSA3.2	CSC	18
WP6-G	TSA3.2	CSIC	45
WP6-G	TSA3.2	CYFRONET	6
WP6-G	TSA3.2	EMBL	5
WP6-G	TSA3.2	INFN	36
WP6-G	TSA3.2	TCD (ended 31-12-12)	17
WP6-G	TSA3.2	UI SAV (shift 10 to TSA2.6)	8
WP6-G	TSA3.3	INFN (shift 24 to CERN, 12 to SA2.6)	24
WP6-G	TSA3.3	CERN	251
WP6-G	TSA3.4	CNRS	29
WP6-G	TSA3.4	EMBL (reclaimed 32 PMS)	5
WP6-G	TSA3.5	INFN	30
WP6-G	TSA3.6	KIT-G	27

SA3 Effort

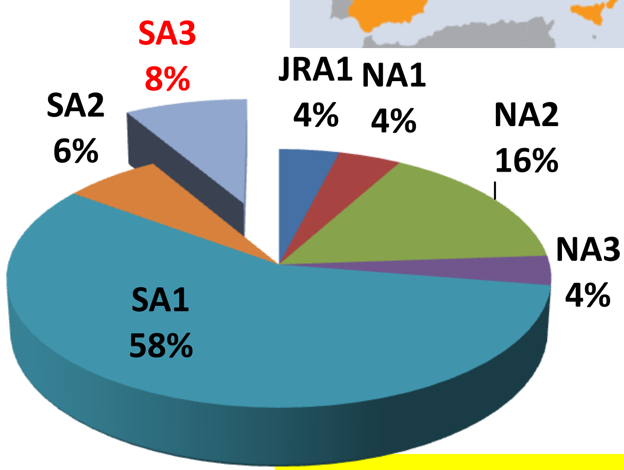

Task ended on 30-04-13 - 95% of total PMs committed have been used

SA3 Objectives

➢ **Transition to sustainable support:**

- + **Identify tools of benefit to multiple communities**
- – Migrate these as part of the core infrastructure
- + **Establish support models for those relevant to individual communities**

Who are the HUCs?

- The communities identified as “Heavy Users” of the grid infrastructure are:
  - – High Energy Physics (HEP);
  - – Life Sciences (LS);
  - – Astronomy & Astrophysics (A&A);
  - – Earth Sciences (ES);
  - – Computational Chemistry and Materials ST;
  - – Fusion.

Achievements in Context

- As a clear example, we use the case of HEP / support for WLCG

✧ The 3 phases of **EGEE** (I/II/III) overlapped almost exactly with **final preparations** for LHC data taking:
  – WLCG Service Challenges 1-4, CCRC'08, STEP'09

➢ **EGI-InSPIRE SA3 covered virtually all the initial data taking run of the LHC: first data taking and discoveries!**
  ➢ **"Mis-aligned" by just over 1 month at the beginning, plus two months at the end**

✓ **The transition from EGEE to EGI was non-disruptive**
✓ **Continuous service improvement has been demonstrated**
✓ **Problems encountered during initial data taking were rapidly solved**
✓ **Significant progress in the identification and delivery of common solutions**
✓ **Active participation in the definition and realization of the future evolution of WLCG**

➢ **All tasks have contributed to adding cohesion to the respective communities (LS, A&A, ES) and preparing for long-term (post-2020) future**

Outlook

- "HUCs" will continue to be (big) users of e-Infrastructures in the future
- A number of major projects are on – or beyond – the (2020) Horizon
  - e.g. Square Kilometre Array, High-Luminosity LHC, EUCLID, etc.
- Will continue to work together, as well as through official bodies such as e-IRG (& User Forum?) to make requirements known

Long-Term Impact

- SA3 has helped increase internal and cross-community cohesion
  - More funding would have improved the latter
- It has contributed to a range of collaborative developments and support
- Specific technical support – e.g. to the LHC programme – would have been missed
  - The recognition at the "Higgs day" might not have happened – justifiably

Post-SA3 Support?

- Future support (funding) could build on successes of SA3, particularly in cross-community support
  - e.g. “Big data” users: Elixir, SKA, HL-LHC
- Period of “meta-stability” is over: we are all facing major platform-related changes (GPUs etc.) and “beyond Moore’s law” needs
- Future e-Infrastructures will still need support: e-IRG message re: User Communities

And that metric?

- To find the Higgs you need 3 things:

1. **The machine;**
2. **The experiments;**
3. **The GRID**

- Rolf-Dieter Heuer, DG, CERN, July 4 2012

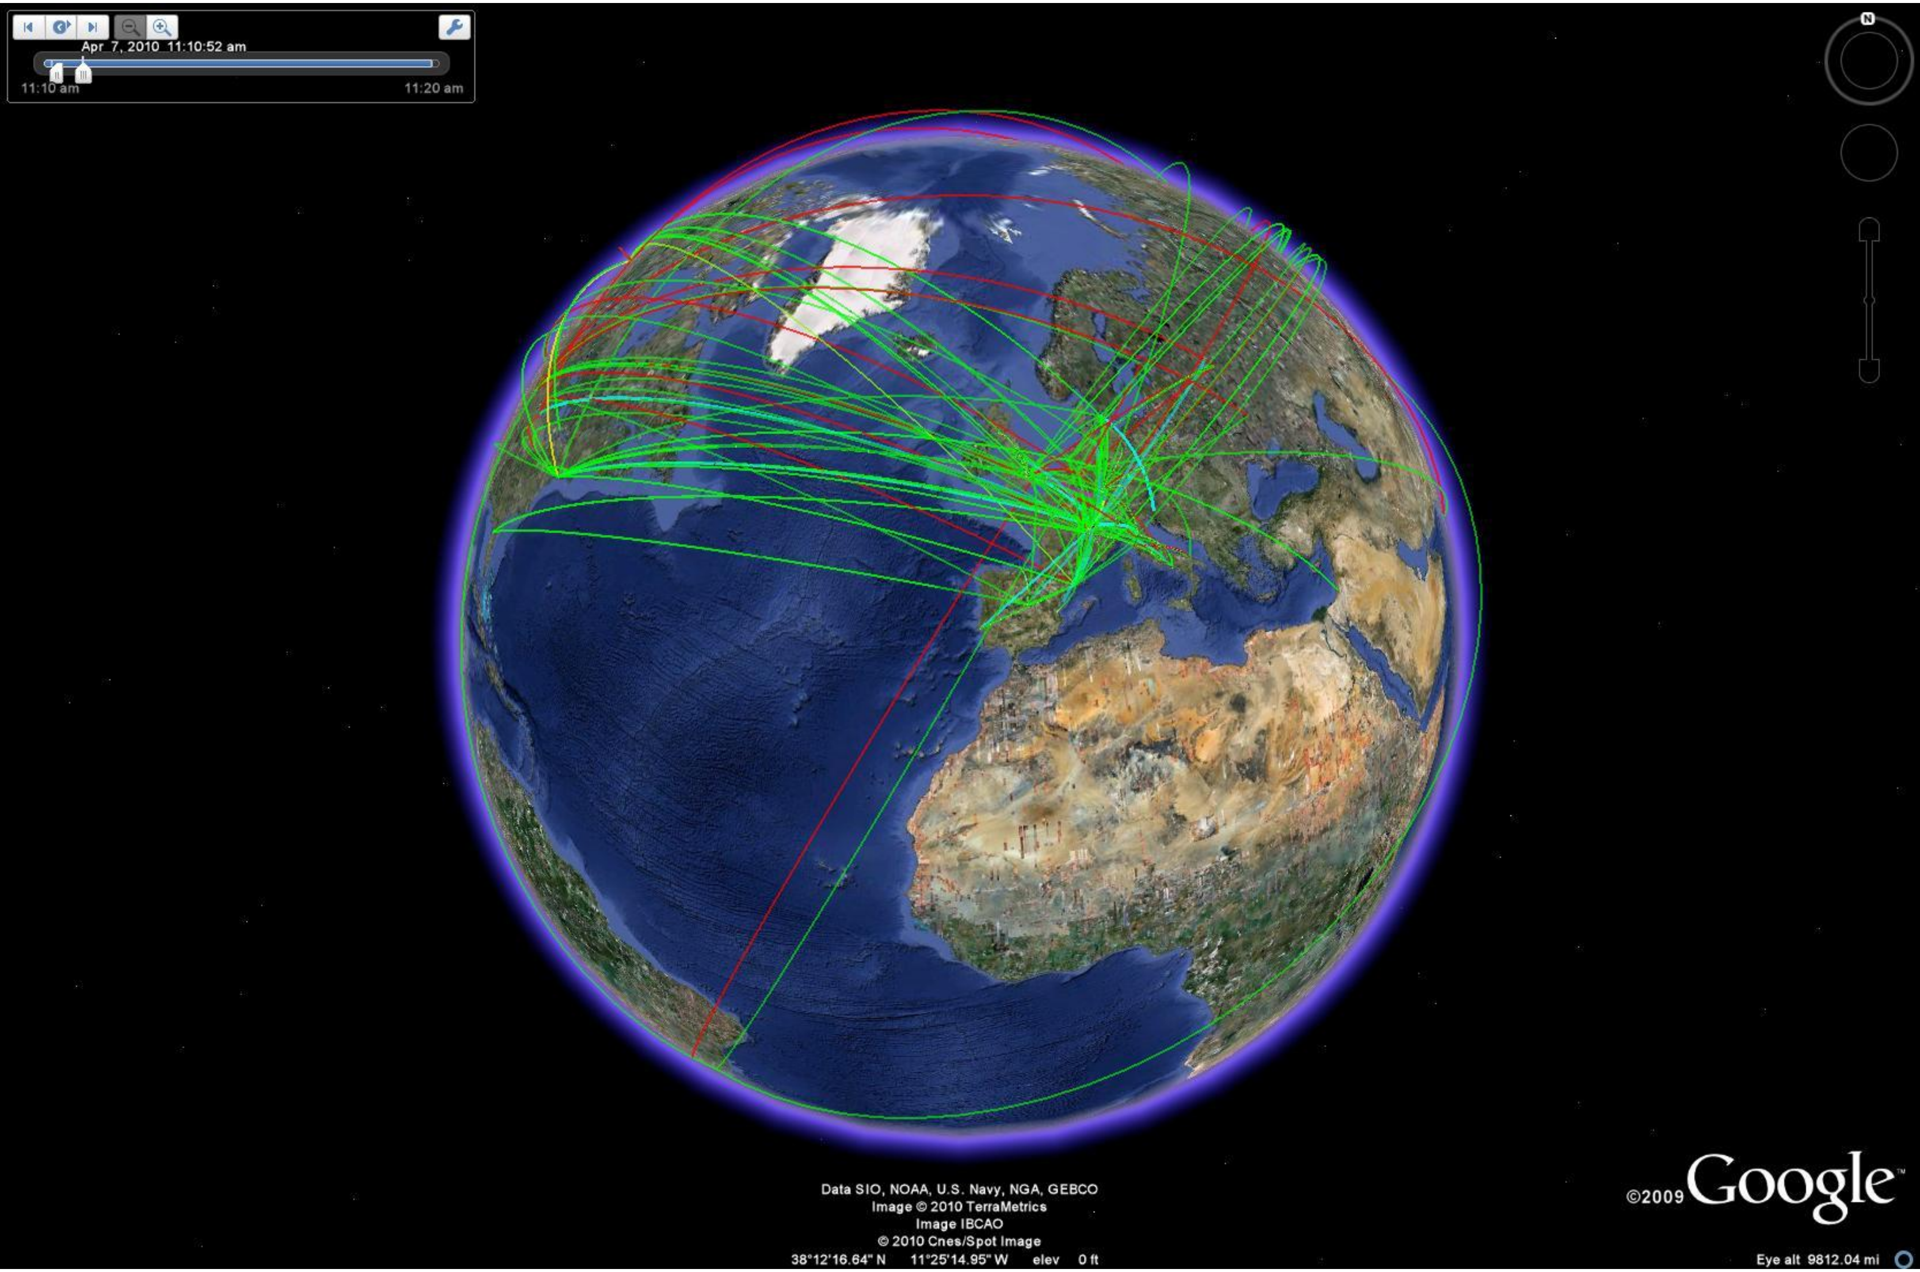
Apr 7, 2010 11:10:52 am
11:10 am
11:20 am
Data SIO, NOAA, U.S. Navy, NGA, GEBCO
Image © 2010 TerraMetrics
Image IBCAO
© 2010 Cnes/Spot Image
38°12'16.64" N 11°25'14.95" W elev 0 ft
©2009 Google
Eye alt 9812.04 mi