

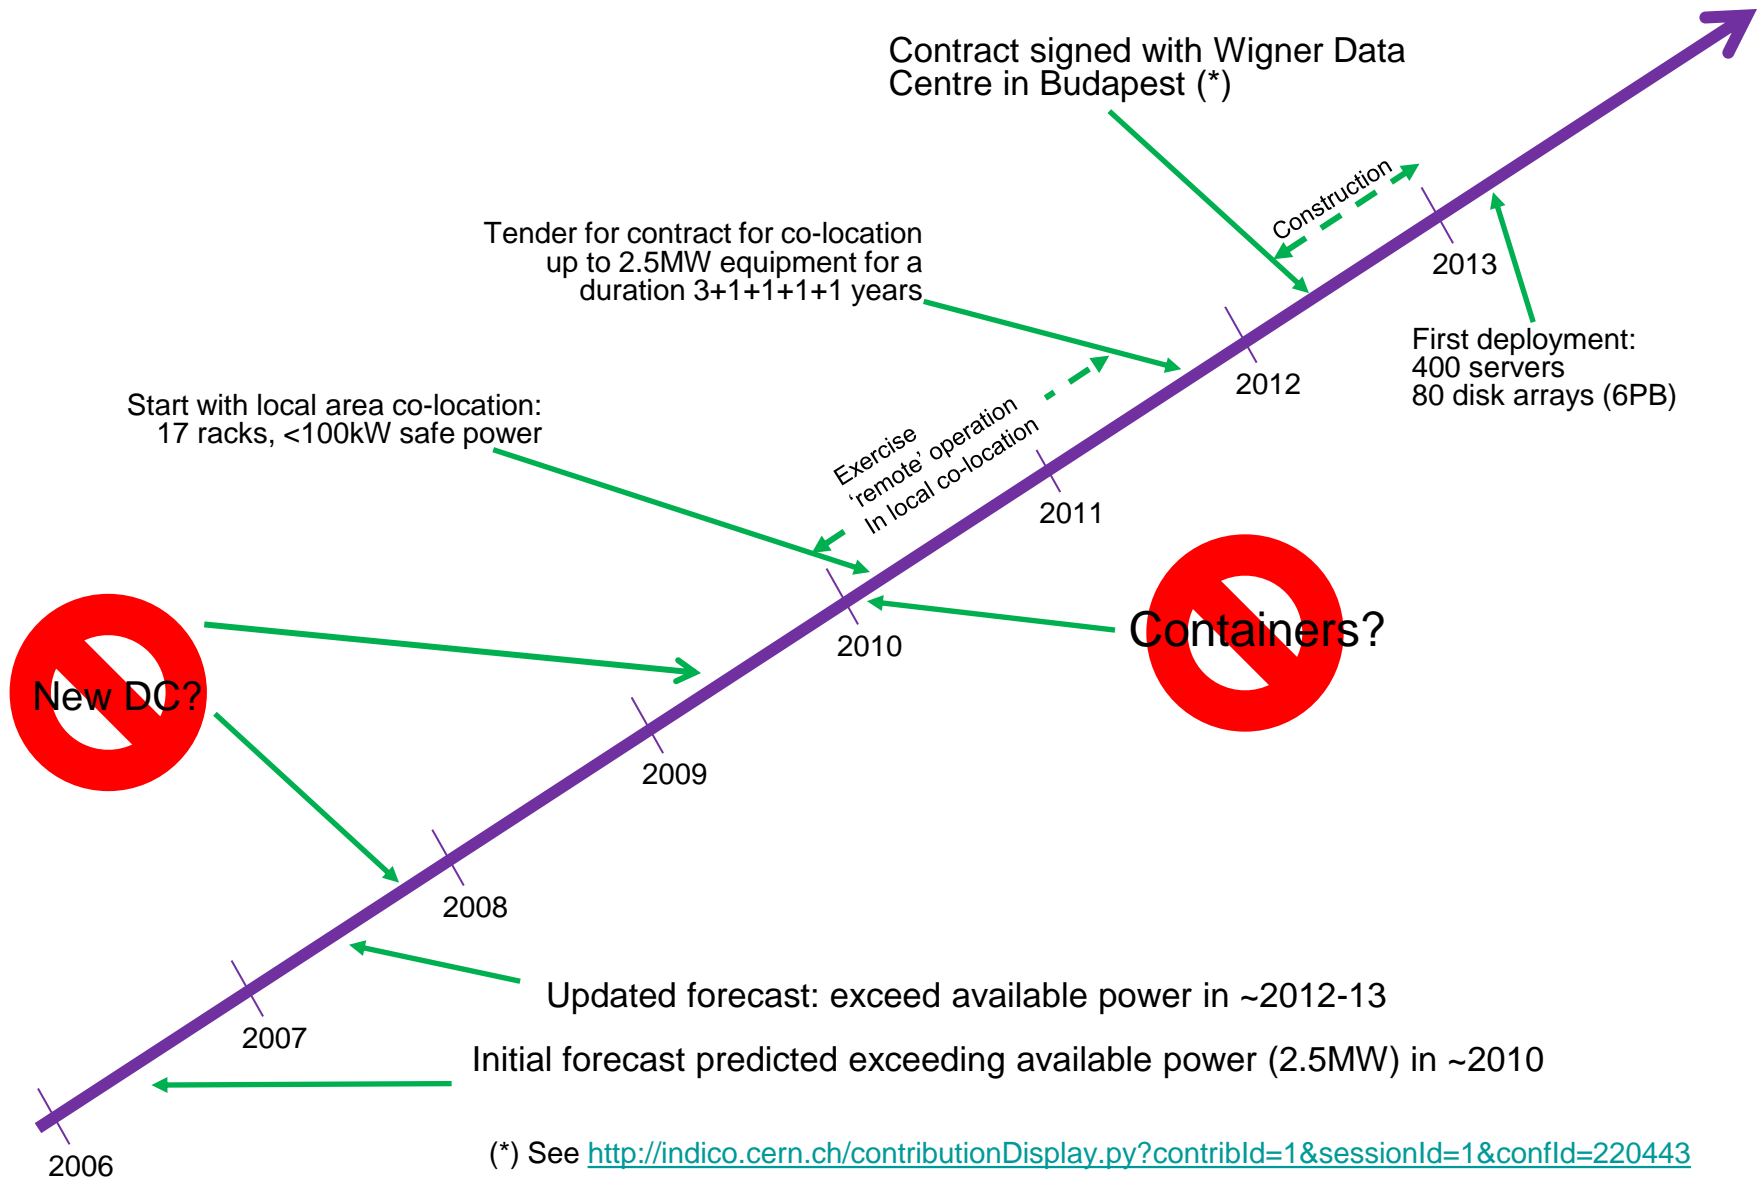
Experience with procuring, deploying and maintaining hardware at remote co-location centre

CHEP'13

14th October 2013

Afroditi XAFI, Alain GENTIT, Anthony GROSSIR, Benoit CLEMENT, Eric BONFILLOU, Liviu VALSAN (since May 2013), Miguel COELHO DOS SANTOS, *Olof BÄRRING*, Vincent DORE, Wayne SALTER

- Background: why remote co-location?
- Preparation
- Hands-on access: why, when, who?
- First remote deployment
- Ramp up remote operation
- Conclusions



(*) See <http://indico.cern.ch/contributionDisplay.py?contribId=1&sessionId=1&confId=220443>

- Review main processes
 - Delivery requirements
 - Hardware handling
 - Stock management
 - Inventory
 - Network registration
 - Burn-in
 - Production deployment
 - Remote console
 - Onsite maintenance

- Review main processes
 - **Delivery requirements**
 - Hardware handling
 - Stock management
 - Inventory
 - **Network registration**
 - **Burn-in**
 - Production deployment
 - Remote console
 - **Onsite maintenance**

- Delivery requirements stipulated in RFP spec and purchase order:
 - Firmware versions & settings
 - Labeling stickers (s/n, MAC, IO ports, disk,, ...)
- Wrong settings tend to break procedures and automation
 - boot order, NIC with PXE disabled, ...
 - struggle with suppliers
- *Remote console access before delivery*
 - *Check and determine detailed settings for supplier*
- *Custom barcode stickers*

- Custom **Asset Identifier** set by supplier

- FRU attributes in BMC

- Contract number in 'Product Asset Tag' (PAT)
- Serial number in 'Product Serial' (PS)

- On chassis

- Bar-code sticker "PAT-PS"



- Network registration:

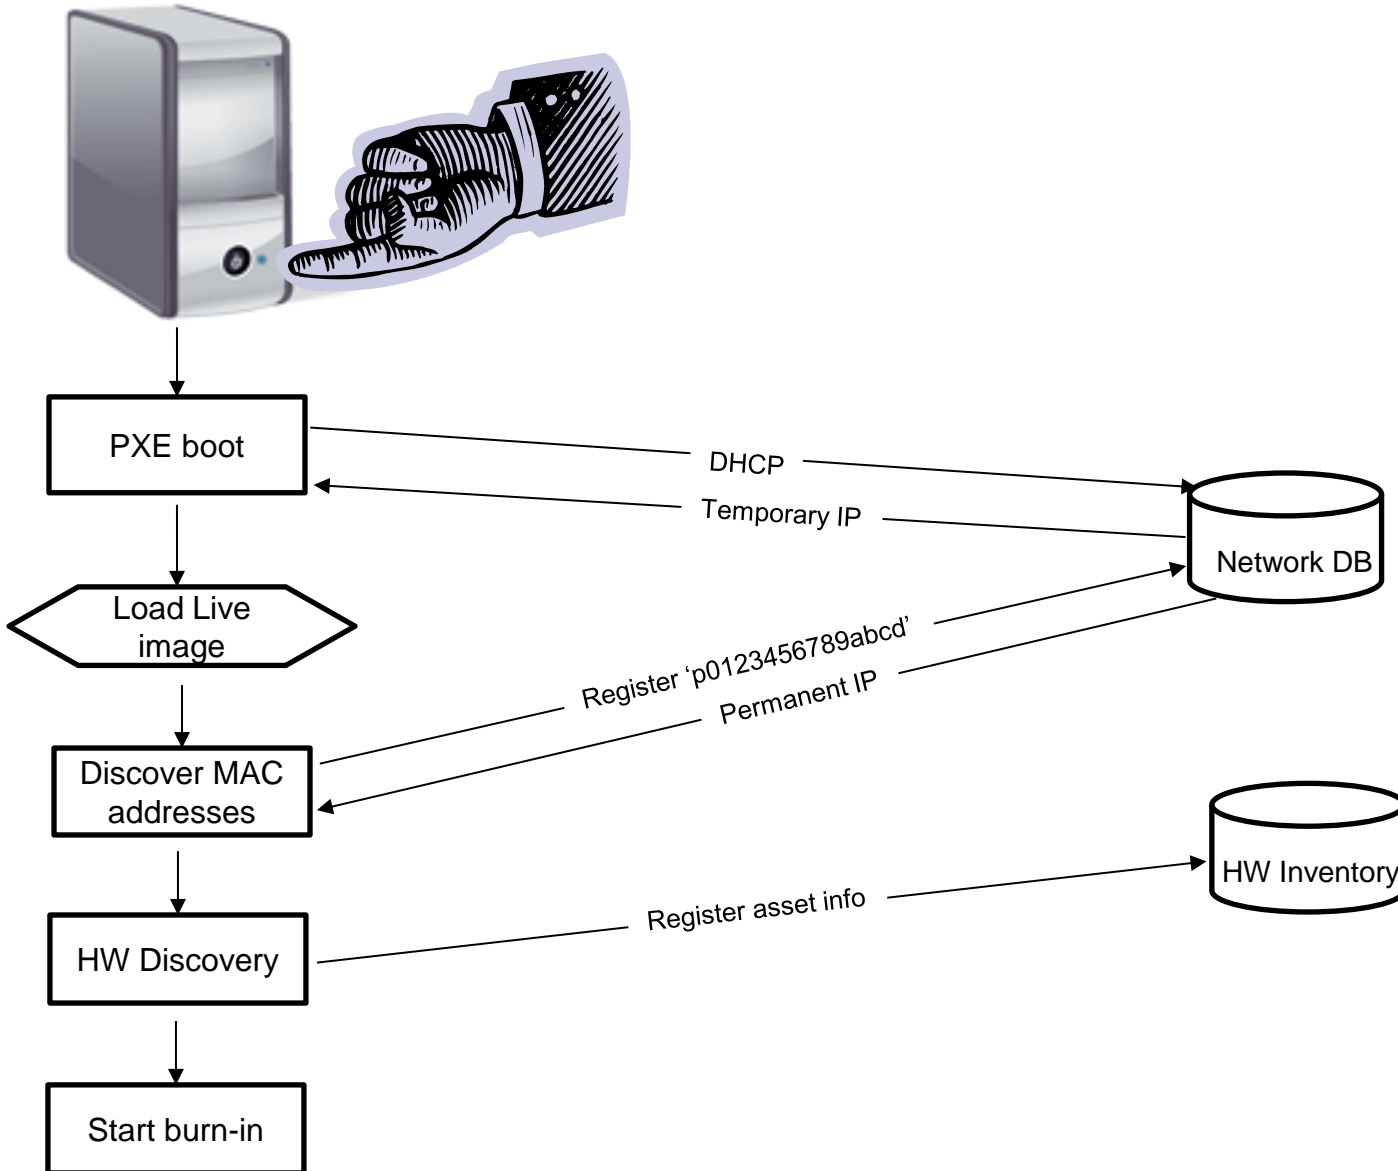
- Host generates its name from Asset ID in BMC

1. **Asset ID too long for Windows NETBIOS name. Compromise:**


Position	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Example	'P'	'0'	'9'	'4'	'7'	'2'	'9'	'6'	'4'	'7'	'5'	'3'	'2'	'7'	'9'
	'P'	Contract CERN doc number							Random decimal number						

2. Host name – Asset id association is stored in the network registration database (LANDB)

Automated registration

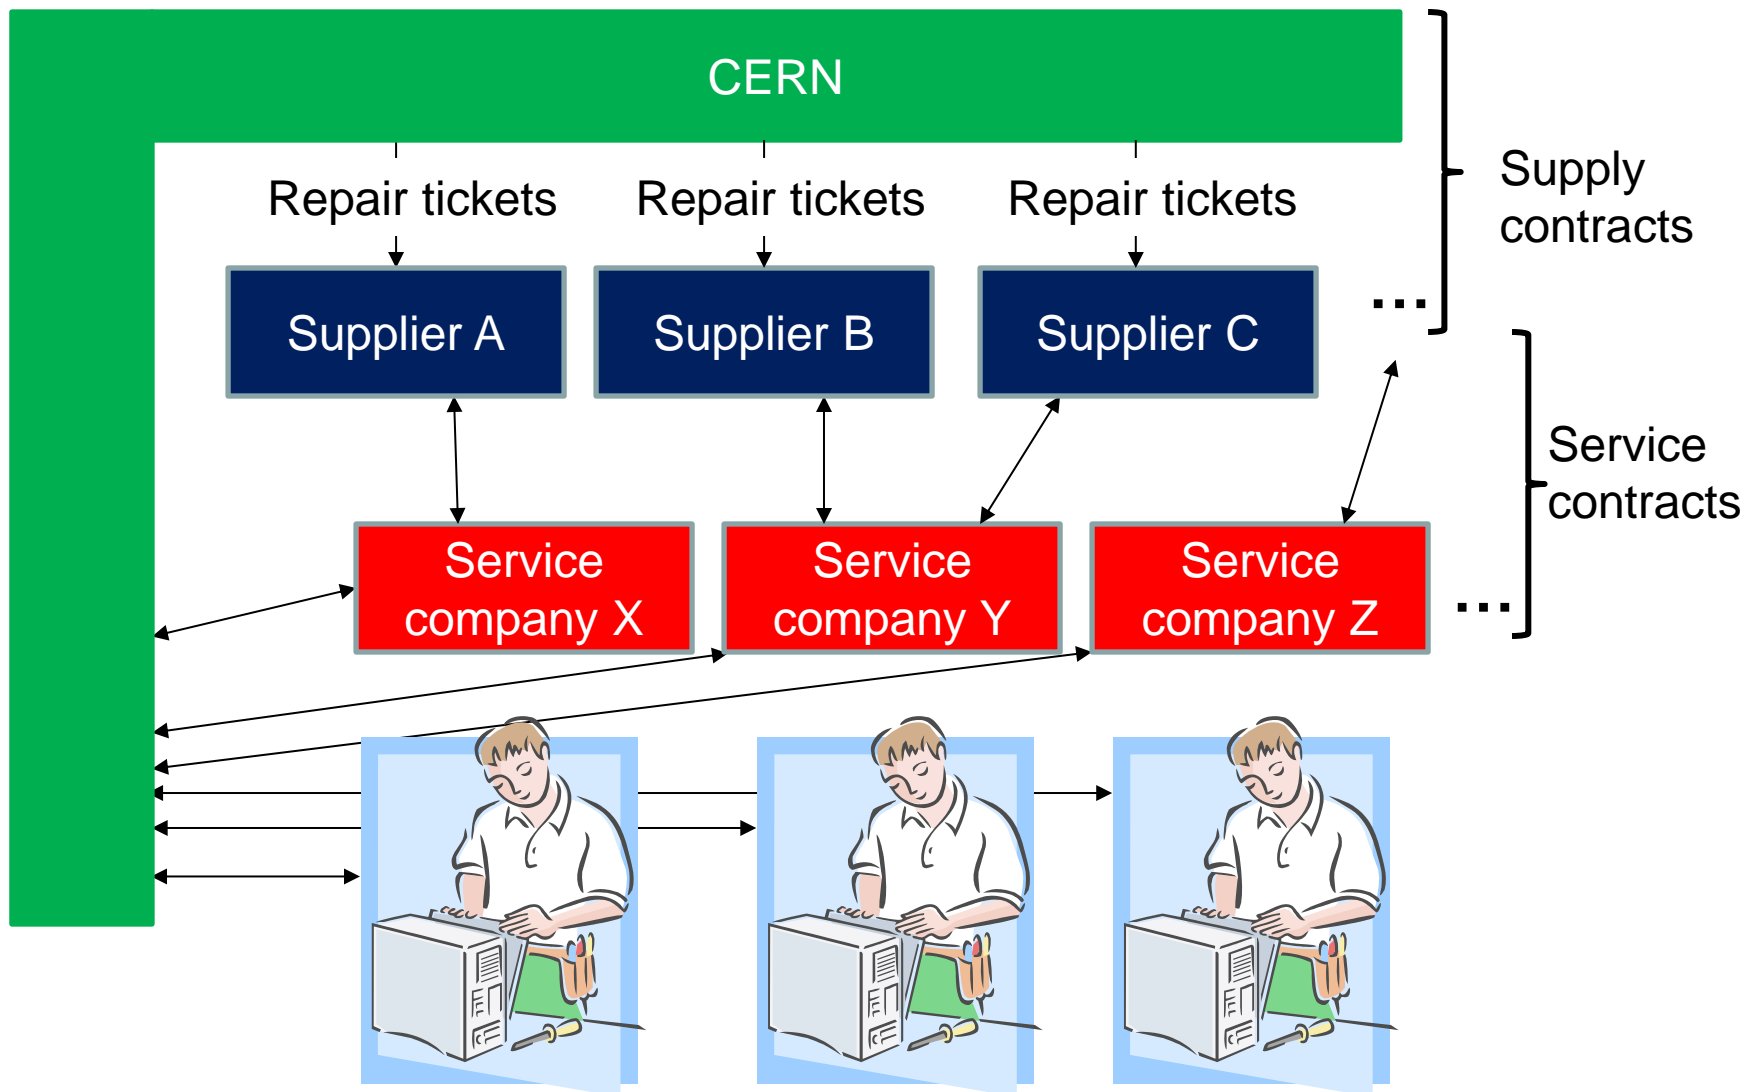


- Runs as part of the live (in memory) image
 1. Memory (memtest) and CPU (burnK7 or burnP6, and burn MMX) endurance tests
 2. Disks endurance tests (badblocks)
 3. CPU and Disk performance tests (HEP-SPEC06, FIO)
 - Network endurance & performance tests (netperf) currently require manual start-up

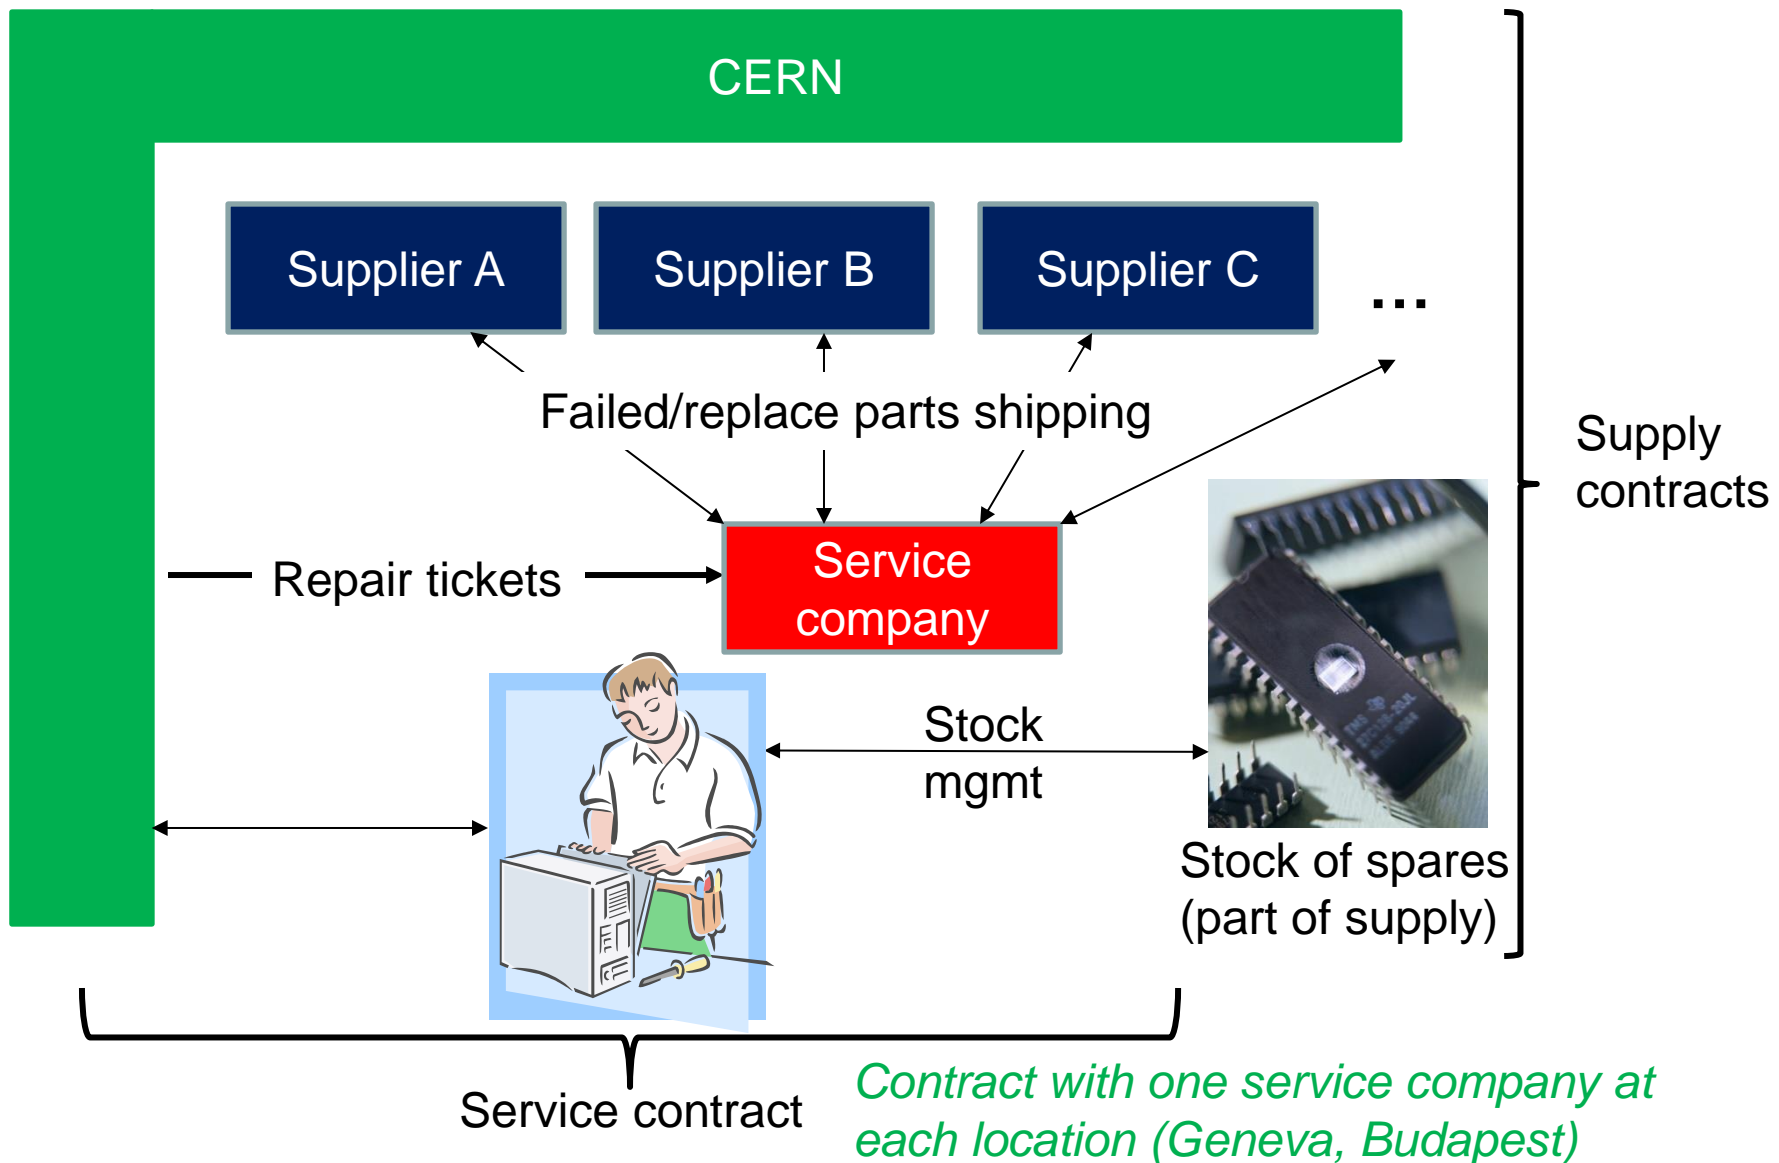


- Guiding principles:
 - If some process can be fully described in a manual procedure...
 - ...it might also be scripted
 - Not always worthwhile in short term
 - Resilience is paramount
 - Failures unavoidable and usual require manual action
 - BUT, it might be possible to carry-on anyway

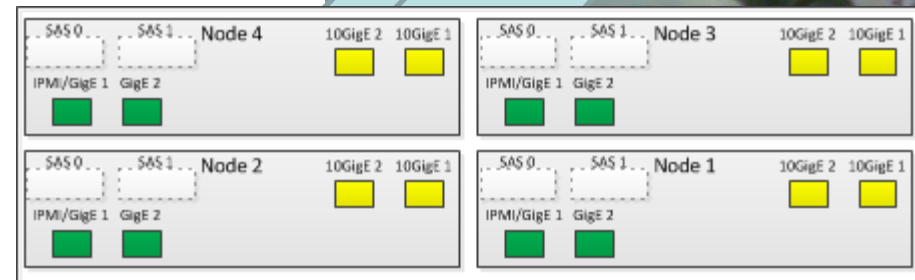

Onsite maintenance <2012



Onsite maintenance ≥ 2012



- Why, When, Who?
 - In principle only to rack mount, cable and repair
- But...
 - Remote console missing or not enough
 - BMC stuck or remote access not working
 - Use of switched PDUs helps
 - Wrong settings
 - improve delivery process!
 - Cabling
 - **Unavoidable!**
 - Risk mitigation:
 - Label with i/o ports
 - Cabling diagrams
 - Resilient automation
- Restricting physical access
 - **Unpopular but unclear why...**
 - ... say hello?





- In autumn 2012 we sent out two RFPs
 - Servers: 300'000 HEP-SPEC06
 - Storage: 28PB raw disk in JBODs
- Delivery to CERN and Wigner
- Two supply contracts per tender
 - Deliver 35% of servers to Wigner
 - Deliver 25% of storage to Wigner
 - + stock of spare parts for on-site repairs

Power on



Registration + burn-in



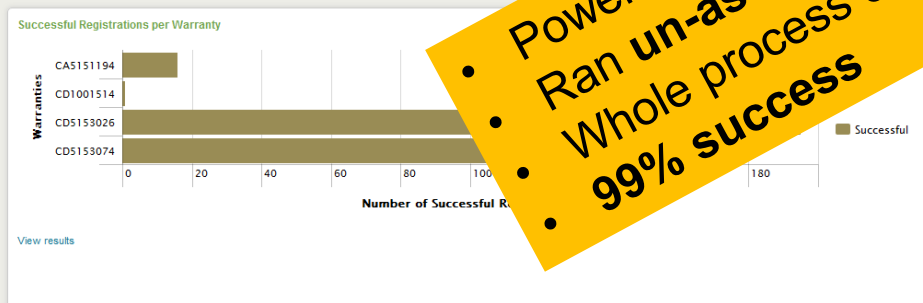
Still in progress

Warranty:

103 Events in the last 7 days (from 16:01:48 April 3, 2013)

uid	hostname	warranty	serial	hostip
9597af...	ca5151876-6n004439ts-2	CA5151876	6N004439TS-2	10.67.7.178
	ca5151876-6n004436ts-2	CA5151876	6N004436TS-2	10.67.7.107
	ca5151876-6n004441ts-1	CA5151876	6N004441TS-1	10.67.6.100
	ca5153026-cr018046c	CD5153026	CR018046C	10.41.212.228
	ca5151876-6n004434ts-4	CA5151876	6N004434TS-4	10.67.4.182
	ca5153026-cr018026a	CD5153026	CR018026A	10.41.212.231
	cd5153026-cr018043c	CD5153026	CR018043C	10.41.213.15
	ca5151876-6n004429ts-1	CA5151876	6N004429TS-1	10.67.5.197
	ca5151876-6n004436ts-4	CA5151876	6N004436TS-4	10.67.6.170
	cd5153026-cr018047c	CD5153026	CR018047C	10.41.213.21

• Power up 400 servers + 80 JBODs
• Ran un-assisted
• Whole process completed in 2 weeks
• 99% success



Registrations per Warranty

real-time

No results found. Inspect...

View results

Detailed Successful Registration List

Warranty:

« prev 1 2 3 4 5 6 7 8 9 10 next »

_time	hostname	warranty	serial	status	hostip	args
03/04/2013 16:00:53.837	cd5153026-cr018030a	CD5153026	CR018030A	OK	10.41.212.172	((DeviceName: 'cd5153026-cr018030a', 'ResponsiblePerson: ('Department: 'IT', 'C
03/04/2013 16:00:45.838	cd5153026-cr018062c	CD5153026	CR018062C	OK	10.41.212.167	((DeviceName: 'cd5153026-cr018062c', 'ResponsiblePerson: ('Department: 'IT', 'C
03/04/2013 16:00:45.838	cd5153026-cr018062c	CD5153026	CR018062C	OK	10.41.212.167	((DeviceName: 'cd5153026-cr018062c', 'ResponsiblePerson: ('Department: 'IT', 'C
03/04/2013 15:41:02.293	cd5153026-cr018049d	CD5153026	CR018049D	OK	10.41.213.52	((DeviceName: 'cd5153026-cr018049d', 'ResponsiblePerson: ('Department: 'IT', 'C
03/04/2013 15:40:35.861	cd1001514-hpc-2218910	CD1001514	HPC-2218910	OK	10.41.208.30	((DeviceName: 'cd1001514-hpc-2218910', 'ResponsiblePerson: ('Department: 'IT', 'C
03/04/2013 15:36:24.225	cd5153026-cr018049c	CD5153026	CR018049C	OK	10.41.213.58	((DeviceName: 'cd5153026-cr018049c', 'ResponsiblePerson: ('Department: 'IT', 'C

Detail Failed Registration List

Warranty:

No results found. Inspect...

View results

- Hardware handling
 - Delivery notification
 - VAT exemption
 - Goods reception
 - Scan bar codes
 - Inventory
 - Rack mounting



- Hardware repair
 - Notification tickets (Service Now)
 - Training and documentation
 - Scheduling
 - Stock management (Infor EAM)



Starting now

- Remote co-location is our way to scale beyond local power limitation
- Wigner contract awarded following competitive tender
- Preparation had positive impact also on local operation
 - Design workflows and automation with remote operation in mind
- Production service is up and running
 - But work still required to finalise operational procedures
- Started preparations for large scale (90%) deployment of new deliveries in 2014-15