Wider Support and Skills

Ste Jones, Alessandra Forti

A hierarchy of responsibilities

- Ideas from Jeremy Coles: We have overlapping skillset needs which sit in a hierarchy. The base represents the time needed for each area. Items higher up the pyramid depend on a knowledge of those below.

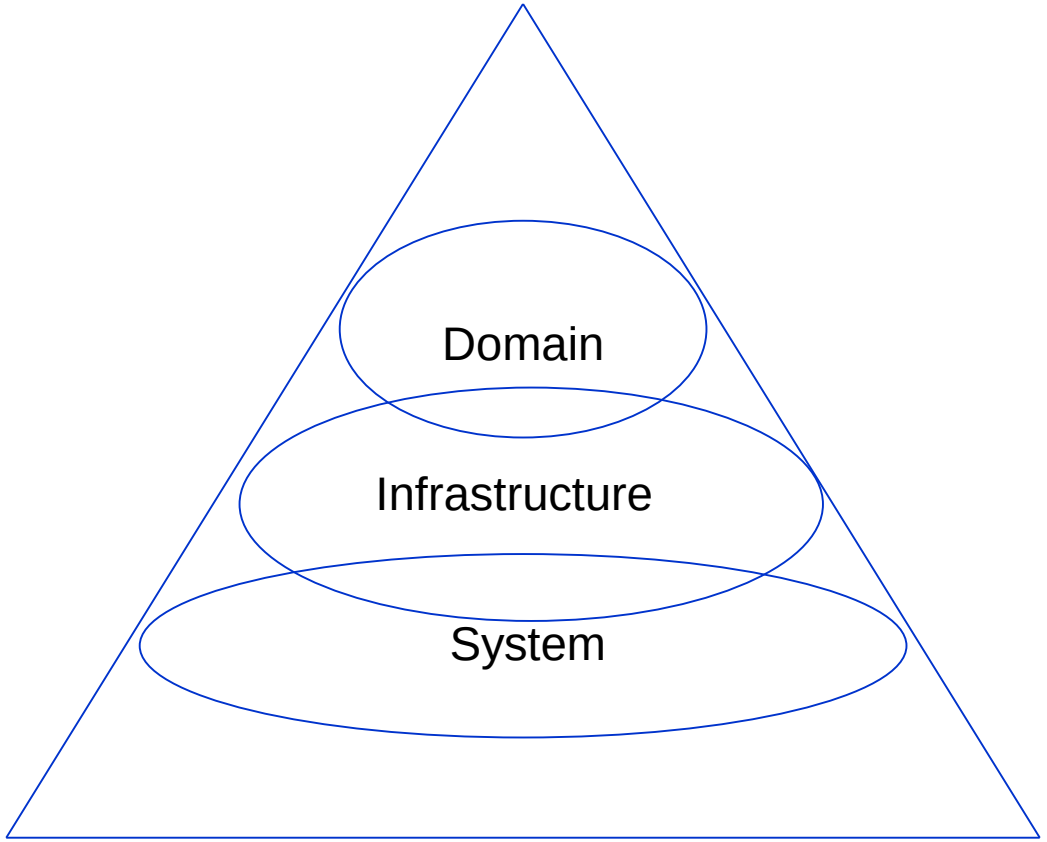

Requires knowledge of:	Example daily tasks:	Additional activities
Experiment workflows;	Support users with job submission;	Study and support performance improvements;
Distributed computing (grid; cloud; federated storage);	Run site/national service nodes;	Test new releases; experiment with configuration improvements;
Compute hardware; machine room dynamics; networking; Operating systems	Install new hardware; patch worker node OS; run monitoring	Tender for hardware; upgrade machine room; test new server types;

Umbrella Concerns

- This talk is based on my experience as a developer and (more recently) systems administrator. There is a set of concerns for us all, whatever our leaning. It think this is a good set to start with; everybody needs to deal with these things.
  - Problem analysis and resolution
  - Security
  - Quality
  - Project management and timelines
  - Capacity analysis
  - Monitoring, benchmarks and tuning
  - Purchasing and other management functions
  - Operational best practises
  - Communications and documentation
  - Conferences, talks, papers and posters
- These concerns, and perhaps others, crop up everywhere, all the time. Furthermore, in my view, any dev, PM or sysadmin, whatever needs to be, first of all, fluent with (say) “Running Linux”, 5th Edition (Matt Welch, ISBN 9784873111315).

Deeper Classification

- Much our work goes beyond traditional Unix Systems Administration. More specialised skills are used in this domain of HEP computing. Some examples of various specialisations might be:
- Research Infrastructure Engineer
  - a leaning towards fabric provision (cpu, storage, network).
- Research S/W Engineer
  - a leaning towards software internals and/or development.
- Site Reliability Engineer (DevOps)
  - a leaning towards automation, monitoring, reliability.

Deeper Classification

- I hope most of us can identify our own roles in at least one of those classifications. I suspect that many people will be involved with more than one of those areas, and other areas beyond these.
- To delve deeper, I'll list out some tasks that are done. This is a brain dump, i.e. these are only the tasks I'm aware of, so please let me what I've missed. Books have been written about this … so I leave many of the detailed implications of these line items to your imagination.
- As general rules for what follows:
  - A) many of these individual tasks apply to the compute, network and storage work areas.
  - B) the tasks generally involve the initial setup or configuration, followed by a long period of on-going maintenance.
  - C) the work spreads over the areas of operating systems, middleware and science application development.
  - D) All of this work is predicated on original feasibility research and tests.

Tasks

- Hardware systems engineering
  - initial setup and ongoing maintenance of hardware systems for compute and/or storage.
- Software systems engineering
  - initial configuration and ongoing maintenance of software systems for compute and/or storage.
- Network
  - initial configuration and ongoing maintenance of network system.

Tasks

- Research and development
  - Feasibility research on the usefulness of emerging technologies and trends (data lakes, e-infrastructure ecosystem …)
  - Middleware integration research to test and integrate potentially useful emerging or existing technologies into production.
  - The design, code, test, deployment and support/maintenance of software or even hardware systems.

Tasks

- System reliability
  - Build and version control automation
  - Backup, recovery and redundancy of systems.
  - Continuous software and hardware migration at various levels:
    - Operating systems
    - Applications
    - Middleware and support software

Tasks

- User/Experiment Support
  - Initial VO On-boarding
    - This is an essential project management/communication task for all new VOs. It involves at least the following sub-steps.
    - Identify organisation stakeholders.
    - Identify organisation requirements and rationalise.
    - Design an initial baseline for organisation.
    - Obtain support at a subset of UK sites, and enable compute/storage.
    - Training for the organisation (esp. tools, certs, proxies and VO set-up or enrolment steps, job design and control etc.)
    - Small scale tests to check feasibility.
    - Large scale roll-out, and continuing operations.

Tasks

- User/Experiment Support (cont.)
  - Ongoing VO support
    - Continuation of the on-boarding activities.
    - Problem analysis and resolution.
    - Technical queries.
    - On-going training and documentation.
    - Baseline evolution.
    - Capacity negotiation

Tasks - support T2s give to VOs (ongoing/growing)

Recent examples from Alessandra….

LSST/SKA

- Basic submission via region. Certificate management.
- Create functional tests (existing Vos)
- Negotiate with VO authorization infrastructure
- Migrate users to “home” VO
- Assist (complex) job submission failure debugging (e.g. Panda and DIRAC tests)
- Help drive technical choices
- Assist users with data migration (DIRAC quirks, resolve Rucio misconfigurations, support protocol testing)
- Support infrastructure tuning/scaling (e.g. debug WMS issues)
- Monitor and trace jobs until workflows work!
- Respond to wider infrastructure changes (e.g. VOMS migration).

Icecube

- Setting up jobs to run on GPUs

ATLAS

- Setup infrastructure for multicore (also needed for others like LSST/SKA)
- Memory based brokering
- Oversee accounting monitoring/corrections
- Arrange site OS migrations
- Coordinate and test deployment of containers (also benefits SKA due to singularity use)
- Coordinate and test deployment of non-gridftp party copy
- Experiment infrastructure site configuration (e.g. Bham DPM->EOS; ECDF dual storage; Glasgow data centre)
- R&D ahead of service consolidation/removal (e.g. BDII)

Tasks – Services supported (from gpp5)

	Task	Service	UK Requires?	UK Involvement	Funding Source: GridPP FTE	UK Other FTE	EGI to UK FTE	Non-UK FTE	GridPP WP	
EGI Services Service tasks done in the context of EGI (both by the UK and others) for the benefit of the whole international community	1	1st and 2nd level support	Yes	Contributes				2.33		**Summary** The UK delivers about 14% of the effort required to provide the global EGI services as follows: 5.3% funded by EGI; 7.3% by GridPP; and 1.3% from other UK sources.
	2	Acceptance criteria	Yes	Feeds back				0.83		
	3	Accounting and Metric Portal	Yes	None				0.50		
	4	Accounting Repository	Yes	Leads	0.75	0.25	0.50		WP-3	
	5	Catchall services	Yes	None				0.67		
	6	Collaboration tools/IT support	Indirectly	None				1.00		
	7	Incident management helpdesk	Yes	None				1.00		
	8	Message Broker Network	Yes	None				0.42		
	9	Monitoring central services	Yes	Feeds back				1.00	WP-2	
	10	Operations Portal	Yes	None				2.00		
	11	Operations support	Yes	None				0.67		
	12	SAM central services	Yes	None				2.33		
	13	Security coordination	Yes	Co-leads	0.40		0.27	1.00	WP-3	
	14	Security monitoring and related support tools	Yes	Feeds back				0.33		
	15	Service registry (GOCDB)	Yes	Leads	0.25		0.25		WP-3	
	16	Software provisioning infrastructure	Yes	None				1.58		
	17	Staged Rollout Infrastructure	Yes	None				0.83		
WLCG Services Service tasks required at an international level forming contributions to WLCG (all UK contributions required by, or a direct benefit to, the UK)	18	Running Shared Infrastructure Nodes	Yes	Contributes	Total effort contributed by WLCG partners is estimated to be 442 FTE. The UK provides 10.3% of this (8.7% or 38.5 FTE ~~funded from GridPP~~ and 1.5% funded by the Tier-2 institutes)				WP-1 and 2	**Summary** The UK contributes a proportionate share of the effort required to perform the international tasks required to run WLCG
	19	Working with Emerging/New Technology Areas	Yes	Contributes					WP-1	
	20	Experiment Support & Liasion	Yes	Contributes					WP-3	
	21	Management & Operations Planning for WLCG	Yes	Contributes					WP-3	
	22	S/W Verification with Experiment Workflows	Yes	Contributes					WP-2	
	23	Networking, Data Transport & Metrics	Yes	Contributes					WP-2	
	24	Operations Coordination	Yes	Co-chairs					WP-2	
	25	Security policy, operations & coordination	Yes	Co-leads					WP-1 and 4	
	26	Software Support for Experiments	Yes	Contributes					WP-2	
	27	Middleware Staged-Rollout	Yes	Contributes					WP-2	
	28	Tier-0 Functional Services	Yes	Phaseing Out					WP-1	
	29	Tier-1 Functional Services	Yes	Contributes					WP-1	
	30	Tier-2 Functional Services	Yes	Contributes					WP-2	
	31	Enabling Tools and Applications	Indirectly	None						
UK Services Service tasks required at a national level to run the production infrastructure	32	Accounting - Benchmarking & Services	Yes	Core	0.11				WP-3	**Summary** A set of national tasks need to be undertaken by all particpants in WLCG
	33	Certificate Authority CA	Yes	Core		0.5				
	34	Data Support for Experiments & Sites	Yes	Core	1				WP-1 and 2	
	35	Documentation & Training	Yes	Core	0.5				WP-2	
	36	Infrastructure services & testing	Yes	Core	1.6				WP-1 and 2	
	37	Interoperation across e-Infrastructures	Yes	Core	0.2				WP-1 and 2	
	38	Infrastructure Monitoring	Yes	Core	1.5				WP-2	
	39	Non-LHC experiment support	Yes	Core	1.2				WP-2	
	40	'On-Duty' Operations Oversight	Yes	Core	0.3	0.1			WP-1 and 2	
	41	Coordination & Management	Yes	Core	0.6				WP-3	
	42	Running Regional Tools	Yes	Core	0.65				WP-2	
	43	Security Evaluations & Response	Yes	Core	1.25				WP-1 and 2	
	44	Support for Tier-3s	Yes	Core	0.2				WP-2	
	45	Ticket Follow-up	Yes	Core	0.19	0.01			WP-2	

We still have these services and tasks and:

- The technology landscape is evolving (positive)
- Bringing in new communities compounds workloads.

UK Services Service tasks required at a national level to run the production infrastructure	32	Accounting - Benchmarking & Services
	33	Certificate Authority CA
	34	Data Support for Experiments & Sites
	35	Documentation & Training
	36	Infrastructure services & testing
	37	Interoperation across e-Infrastructures
	38	Infrastructure Monitoring
	39	Non-LHC experiment support
	40	'On-Duty' Operations Oversight
	41	Coordination & Management
	42	Running Regional Tools
	43	Security Evaluations & Response
	44	Support for Tier-3s
	45	Ticket Follow-up

Skills

- Below I present an off-the-cuff list of skills; it would be rare (but not impossible) get all these skills in a single person.
- The next step might be to map the tasks lists above to the job classifications. We know, for example that some individuals are highly involved with (say) VO on-boarding, or whatever it is. And others are highly involved with software development or in doing research on integrating emerging technologies.
- Slightly different “leanings” may call for different strengths. Unfortunately, this is where I start to sound like a recruiter. Sorry about that.

Skills

- All roles
  - The background skills for all these rolls are listed above on slide 2. In addition, general useful skills should include quality assurance/CMM, general IT, science or engineering skills, knowledge of Open Source culture, general knowledge of tools (yum, rpms, puppet, ansible), the soft skills, and an inquiring mind. For more specialised roles:
- Research Infrastructure Engineer
  - This involves fabric provision. Therefore, this role would call for particularly solid design, installation and configuration expertise and knowledge in hardware (and on-going maintenance), as well as advanced knowledge of networks and storage technologies, e.g. RAID6, or ZFS etc.
- Research S/W Engineer
  - This would call for particularly solid programming skills (bash, Perl, Python, Java, C, SQL ..), current concepts (Web Services, Databases, networks), knowledge of the e-infrastructure ecosystem, and strong knowledge of current and emerging Grid frameworks (GSI, storage protocols, Cloud, commercial provision, middleware concepts and software internals.) Massive Physics expertise required for development of modelling algorithms.
- Site Reliability Engineer (DevOps?)
  - This is “what happens when a software engineer is tasked with what used to be called operations." It involves strong debugging and workaround skills, decent knowledge of application and OS internals, and total familiarity with automation tools such as Puppet/Hiera or Ansible. According to Wikipedia, “a site reliability engineer (SRE) will spend up to 50% of their time doing "ops" related work such as issues, on-call, and manual intervention. Since the software system that an SRE oversees is expected to be highly automatic and self-healing, the SRE should spend the other 50% of their time on development tasks such as new features, scaling or automation. The ideal SRE candidate is a highly skilled system administrator with knowledge of code and automation.” I’d add QA, PM and S/W Eng skills to that list.

END

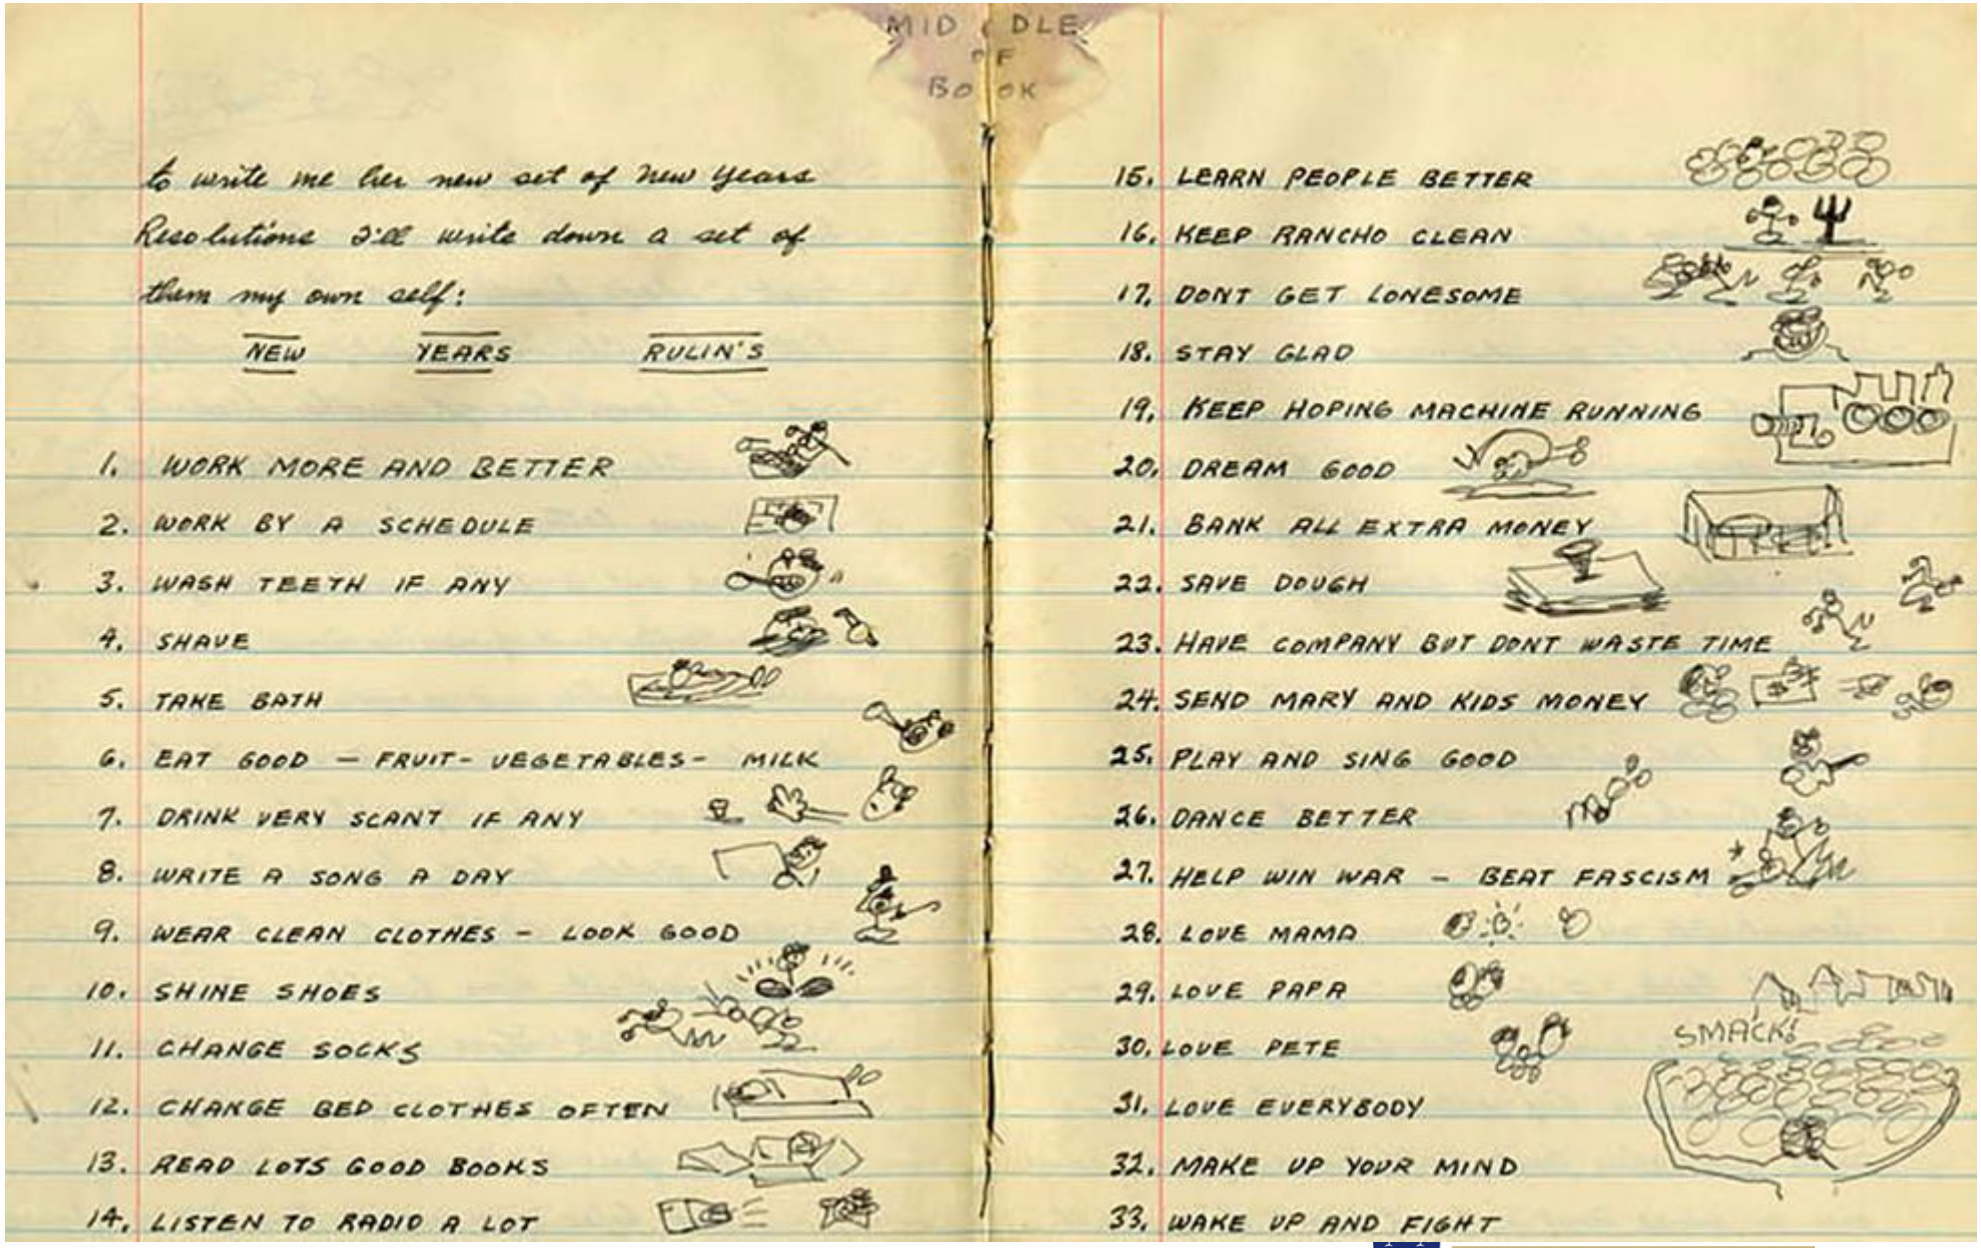

MIDDLE OF BOOK

to write me her new set of New Years Resolutions I'll write down a set of them my own self:

NEW YEARS RULIN'S

1. WORK MORE AND BETTER
2. WORK BY A SCHEDULE
3. WASH TEETH IF ANY
4. SHAVE
5. TAKE BATH
6. EAT GOOD — FRUIT - VEGETABLES - MILK
7. DRINK VERY SCANT IF ANY
8. WRITE A SONG A DAY
9. WEAR CLEAN CLOTHES - LOOK GOOD
10. SHINE SHOES
11. CHANGE SOCKS
12. CHANGE BED CLOTHES OFTEN
13. READ LOTS GOOD BOOKS
14. LISTEN TO RADIO A LOT
15. LEARN PEOPLE BETTER
16. KEEP RANCHO CLEAN
17. DONT GET LONESOME
18. STAY GLAD
19. KEEP HOPING MACHINE RUNNING
20. DREAM GOOD
21. BANK ALL EXTRA MONEY
22. SAVE DOUGH
23. HAVE COMPANY BUT DONT WASTE TIME
24. SEND MARY AND KIDS MONEY
25. PLAY AND SING GOOD
26. DANCE BETTER
27. HELP WIN WAR - BEAT FASCISM
28. LOVE MAMA
29. LOVE PAPA
30. LOVE PETE
31. LOVE EVERYBODY
32. MAKE UP YOUR MIND
33. WAKE UP AND FIGHT

SMACK!