



Government  
of Canada

Gouvernement  
du Canada



# Canadian Meteorological Centre HPC Renewal Initiative

iCAS2015, September 16 2015

Laurent Chardon, Environment Canada

Luc Corbeil, Shared Services Canada

Canada 

# Agenda

---

- Background information on the project
- Environmental Forecasting Requirements
- Supercomputing Procurement
- Conclusion

# Background

---

- Contract signed in 2003 for the Scientific Computing Facility (SCF) with IBM, now soon to expire.
- Additional funding recently awarded to the Meteorological Service of Canada
  - Component 1: Monitoring Networks
  - Component 2: **Supercomputing capacity**
  - Component 3: Weather Warnings and Forecast System
- As a result: HPC Renewal Initiative launched

# Variety of customers of SCF

---

- Weather forecasting
  - Research
  - Development
  - Operations
- Climate
- Air quality
- Environmental emergency response

# Environmental Forecasting Requirements

---

- Science drives the requirements
- User Requirement Document (URD) produced
  - 10 year plan outlining the expected scientific path (EC)
  - Translated in computing needs (EC&SSC)
  - RFP (SSC lead with EC representation)

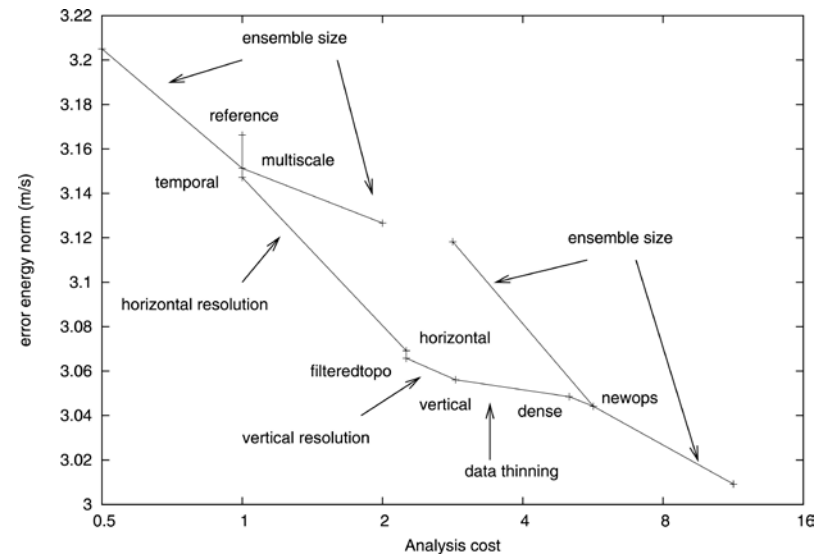
# Meteorological Research Division

---

- RPN-A
  - Atmospheric modelling
- RPN-E
  - Environmental modelling
- RPN-SI
  - Informatics section
- ARMA
  - Data assimilation
- ARMP
  - Cloud physics and severe weather

# Atmospheric Numerical Prevision Research

- HR global forecast 15 km, 4/day
  - 80 levels -> 120
- HR North America LAM 250m
- New 3D microphysics scheme
  - Turbulence
- Ensemble forecast to HR 50km -> 25 km -> 15 km
- Increase radiation calls
- Cloud physics renewal



# Environmental Numerical Prevision Research

---

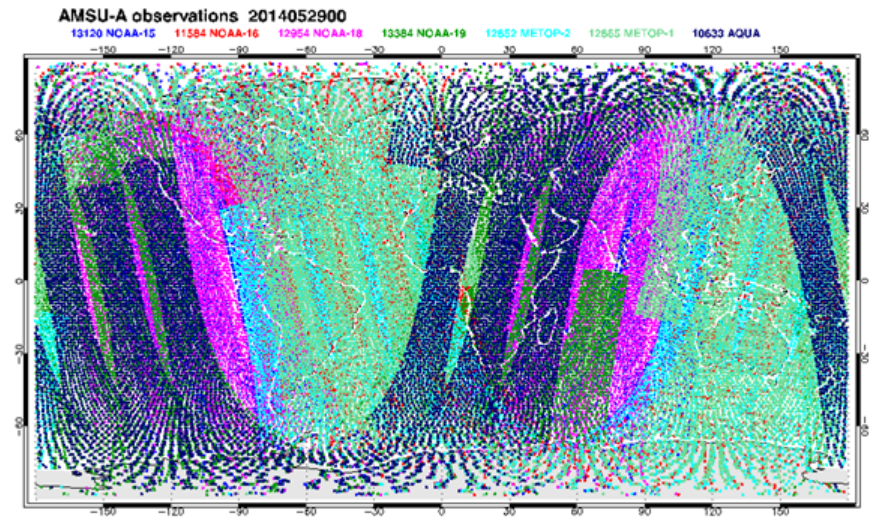
Ocean, ice, waves, urban, ground, vegetation, hydrology, lakes

- Global coupled Ocean/Atm forecast – 15 km
- Global ice forecast
- Global wave forecast
- Ensemble arctic ocean forecast (20 + 20 mercator)
  - 4-5 days
- HR arctic ocean deterministic forecast (1/36 °)
  - < 4 days
- Coupled HR Lake-Ice-Atm over the great lakes



# Data Assimilation

- Data assimilation at same resolution as forecast
- Add aeolus, GOES-R data
- Increase ensemble # of members 256 -> 512
- Surface data assimilation -> soil moisture
  - SMAP
  - O(1000) -> O(15,000,000)



AMSU-A, 6-h period, 7 satellites, temp profiles

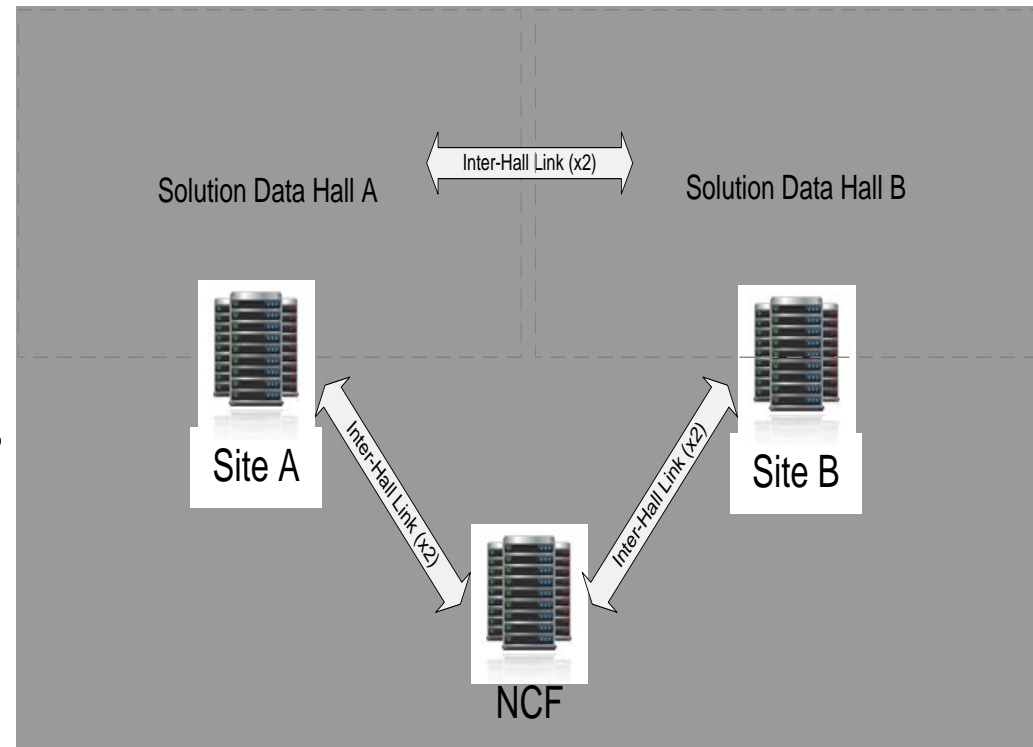
# Shared Services Canada

---

- Created in 2012, to take responsibility of email, networks and data center for the whole Government of Canada.
- Supercomputing IT people working for EC transferred to SSC.
- Scope of the HPC team expanded to all science departments
- As in any reorganization, there are challenges and opportunities!

# EC Supercomputing Procurement

- Contract for Hosted HPC Solution: 8.5 years + one 2.5 year option (three systems + one optional)
- Connectivity between HPC Solution, Data Halls and Dorval
- No more than 70km between Hall A, Hall B & Dorval
- Flexible Options for additional GC needs

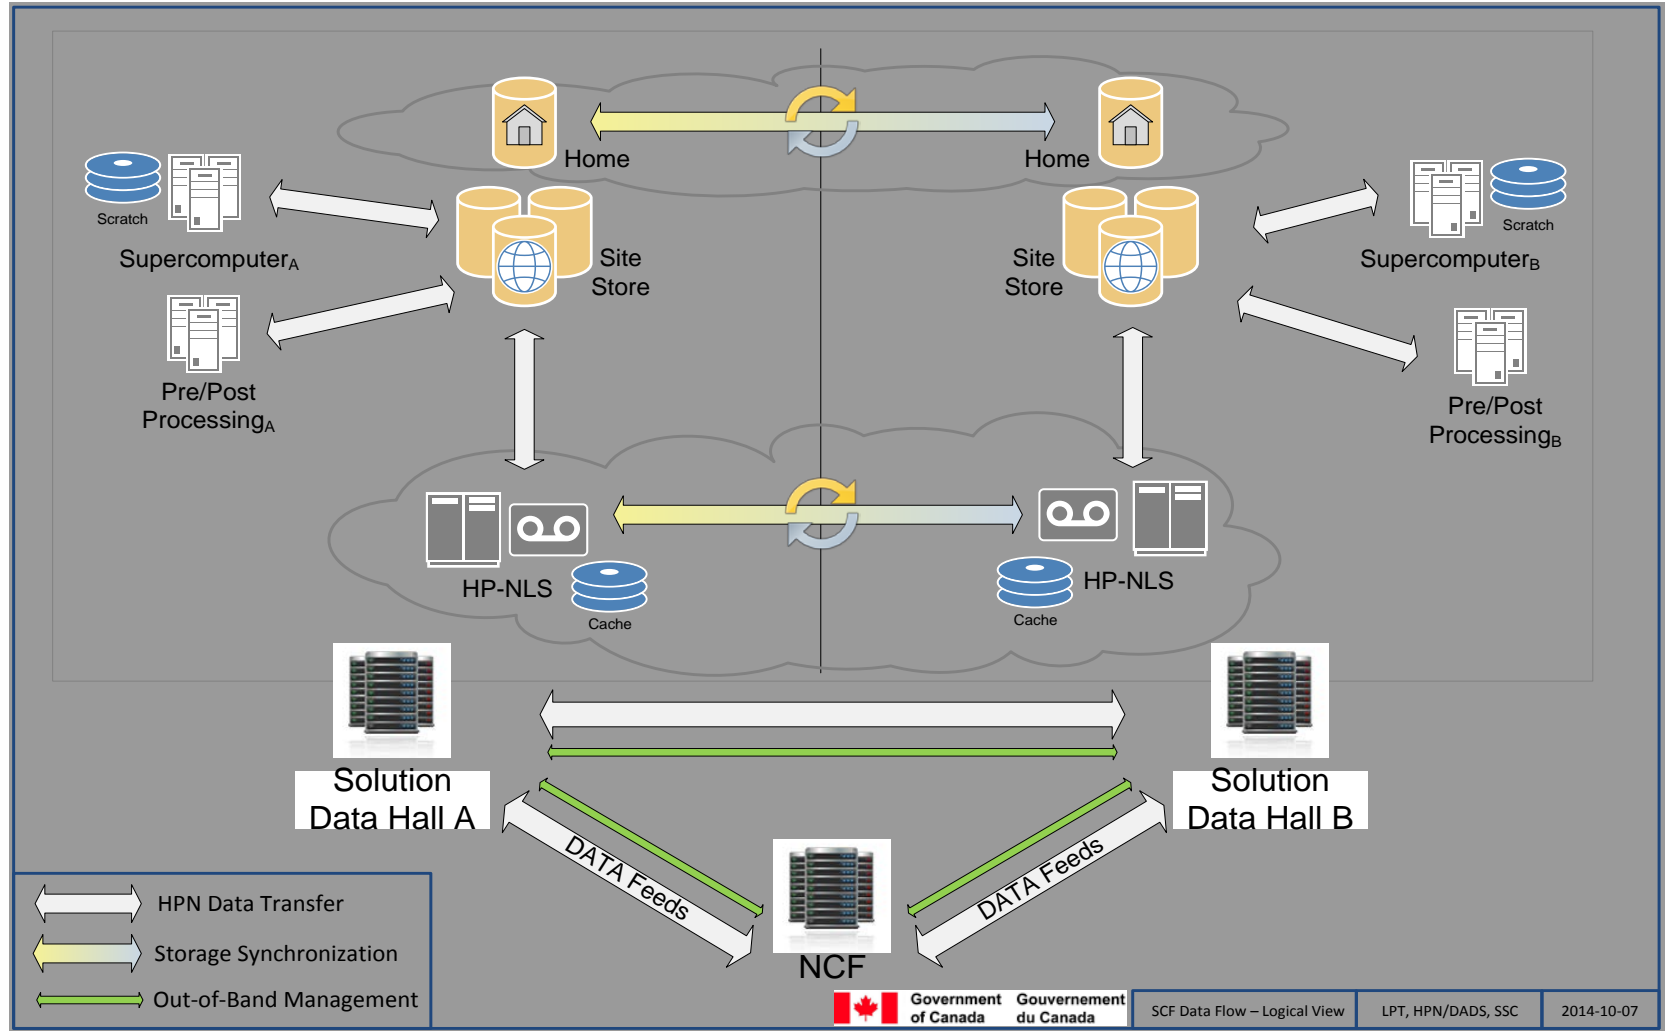


# Scope

---

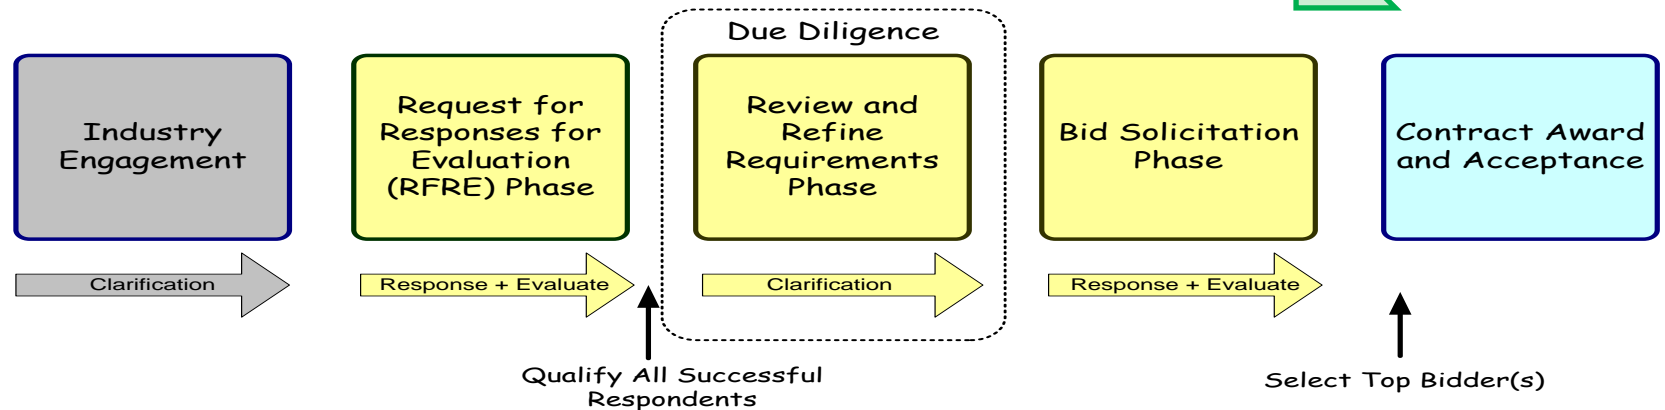
- Supercomputer
- Pre/Post-processing
- Global Parallel Storage (high-bandwidth low latency)
- Near-line Storage
- High Performance Interconnect
- Software & tools
- Maintenance & Support
- Training & Conversion support
- On-going Availability

# HPC Solution: Fully Redundant



# Collaborative Procurement Process

- Use of (more) Collaborative Procurement Process
  - ✓ RFI
  - ✓ Invitation to Qualify
    - 4 Qualified Respondents (QR): Cray, Dell, HP & IBM
  - ✓ Review and Refine Requirements Phase
  - ✓ Bid Solicitation Phase
  - Contract Award



# Technology

---

- Quite a bit of freedom on how the vendor can achieve the performance targets
- A lot of restrictions came from applications that aren't in the benchmark suite. They need to perform as well
- No GPUs nor other accelerators on 1<sup>st</sup> system, strong possibility on the 2<sup>nd</sup>

# Bid Evaluation

---

- One number that drives almost everything: Fixed Performance Level (FPL)
  - From that number, we derived floor amounts for memory, storage, benchmark performance, tape qty, etc.
  - Minimum value is 0.25 (lower bid == non compliant)
    - It means ~5X performance increase (with FPL= 0.25. ~20X for FPL=1) for 1<sup>st</sup> system on the main GEM benchmark over the Power7, then 2.6X for each upgrade
- Requirements increased automatically as closing date of the RFP got pushed
  - On an exponential growth curve, the starting point matters a lot!



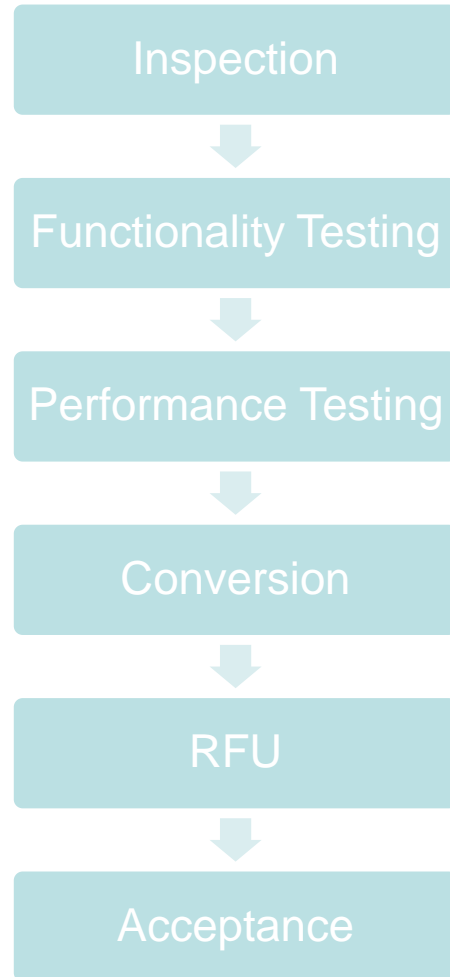
# Reality Check

---

- Tender issued late November
- Supply Chain Integrity (Security Phase)
  - QR's supply chain is vetted by GoC for security purposes
  - Took much longer than anticipated
- Market environment changes
  - Procurement risks, technology, foreign exchange rate, etc.
- And... Federal Elections
  - Longest election campaign ever in the country's history launched in August (78 days).
  - Most main decisional processes are paralyzed until November... at the earliest

# HPC Implementation Milestones: Delivery to Acceptance

- Data Hall and Hosting Site Certification
- Functionality Testing (IT infra)
- Security Accreditation
- Performance testing
- Conversion of Operational codes (Automated Environmental Analysis & Production (AEAPPS))
- Meeting the above triggers a 30 day availability test

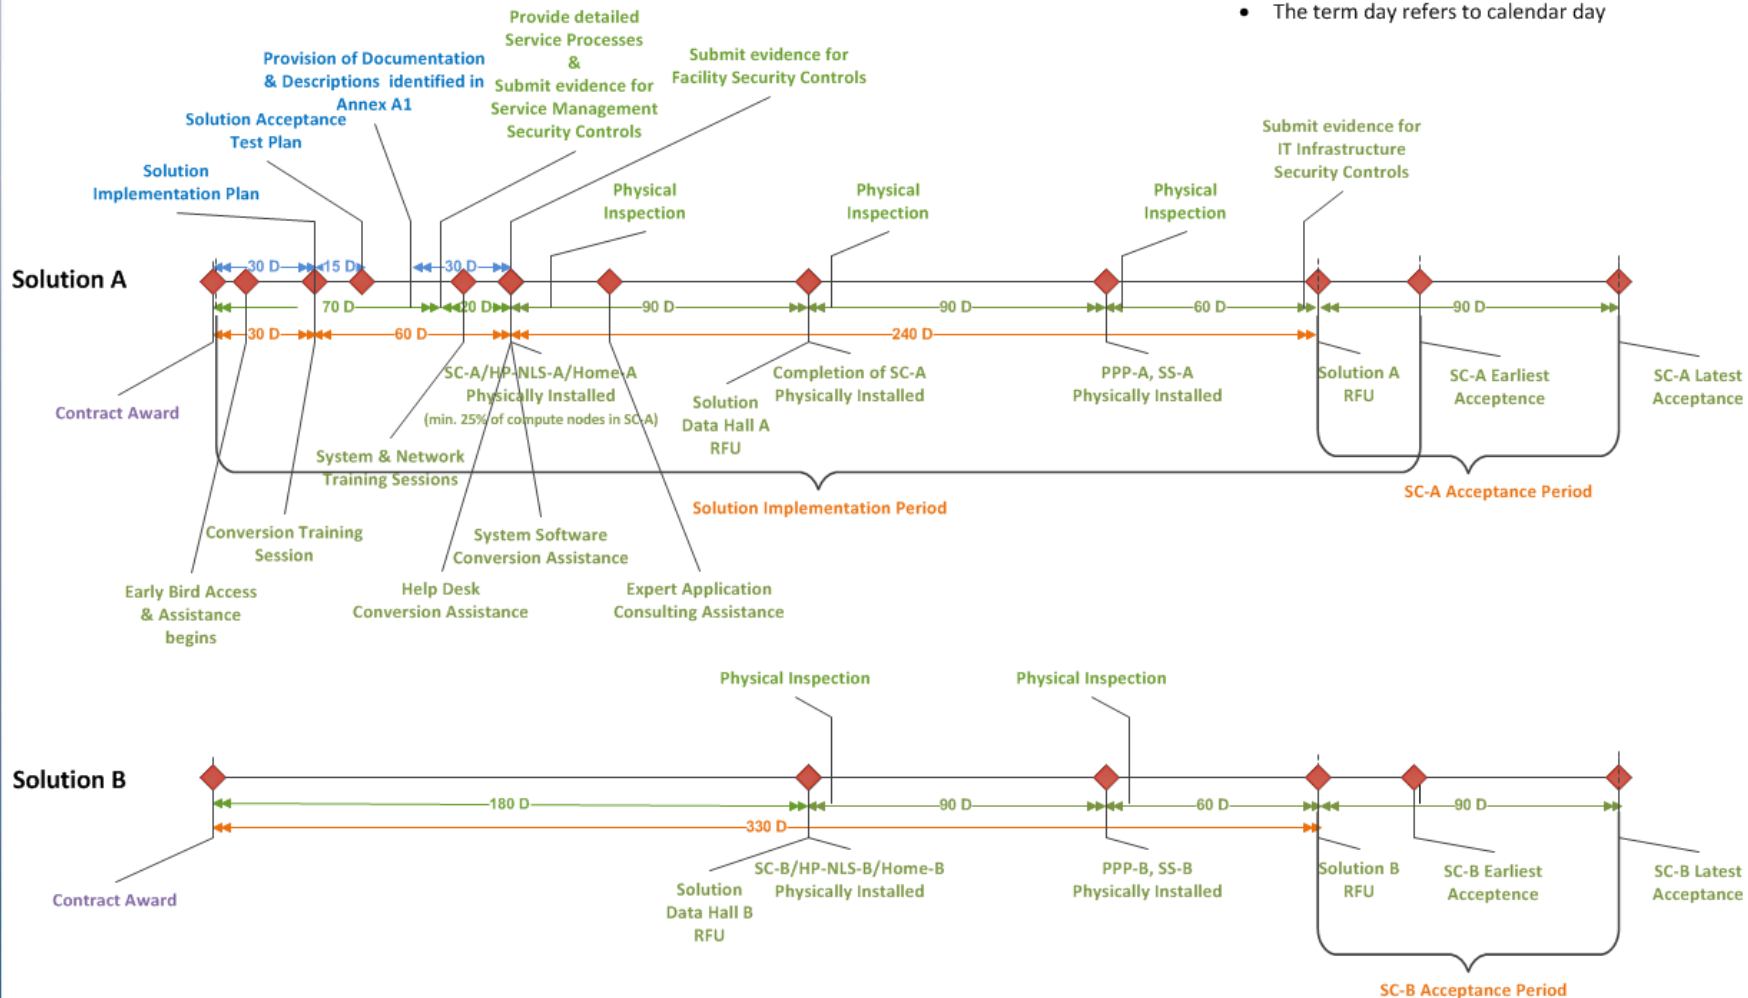


\*Scheduling Details in next slide

# Once the Ink Dries...

## SSC Solution Timeline

- All dates are "no later than", unless otherwise specified.
- The term day refers to calendar day



# General Purpose Science Cluster (GPSC)

---

- Very similar to UK NCAS approach
- Federating users of many computing communities
- Using Linux containers (LXC, not docker)
  - Works well, but not out of the box
- Workloads extremely heterogeneous
  - Month(s)-long jobs?
  - 4TB RAM on the node?
  - 100Mio 32k files? Or one 3.2TB file? Or a 3PB database?
  - Linux and Windows?
- In dev mode. More to report in a year

# Summary

---

- Canada is investing massively in HPC to support the Environmental Prediction program
- RFP almost completed
  - Evaluation of the bids still in progress
- It will eventually bring EC to O(100) Petaflops and O(1) Exabyte