Service | Innovation | Value

Shared Services Canada
Environment and Climate Change Canada
HPC Renewal Project

CUG 2017
Redmond, WA, USA
Deric Sullivan
Alain St-Denis & Luc Corbeil
May 2017

Shared Services Canada | Services partagés Canada

Canada

Background: SSC's HPC Renewal for ECCC

- Environment Canada highly dependent on HPC in delivery of mandate: simulation of Environmental Forecasts for health, safety, security and economic well-being of Canadians.
- 2012: Much of ECCC's IT staff moved over to SSC
- Scope of the SSC HPC team expanded to all science departments
- Joint ECCC-SSC submission for Supercomputing Capacity
- Contract Award (May 27 2016)

Shared Services Canada – Our Mandate

Shared Services Canada was formed to consolidate and streamline the delivery of IT infrastructure services, specifically email, data centre and network services. Our mandate is to do this so that federal organizations and their stakeholders have access to reliable, efficient and secure IT infrastructure services at the best possible value.

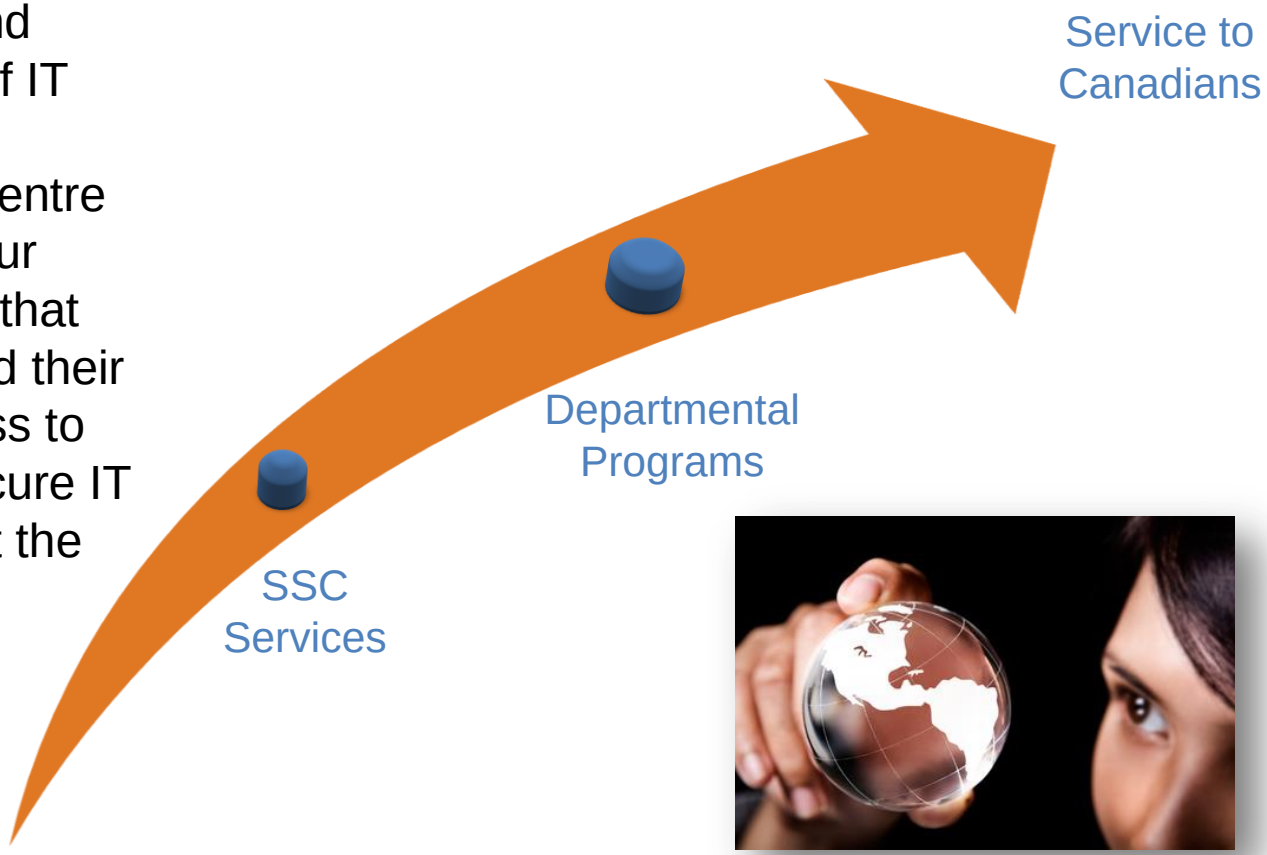

SSC will Innovate, ensure full Value for Money and achieve Service Excellence !

A Bit of History

- ECCC has been using a supercomputer for weather forecasting and atmospheric science for more than half a century

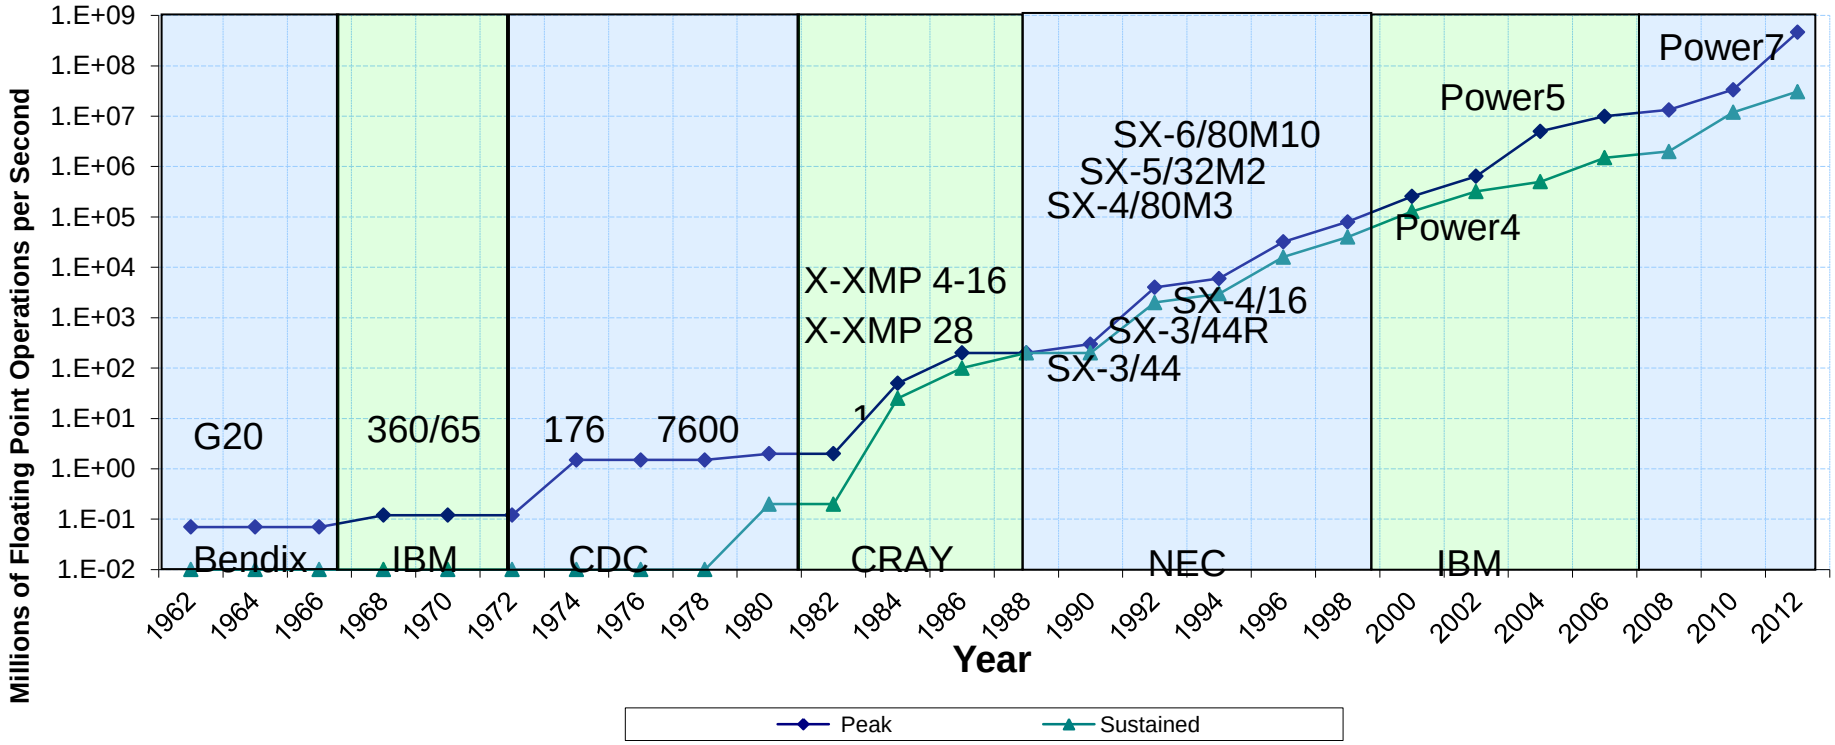

Scope of HPC Renewal

Scope	In replacement of
Supercomputer clusters	Two 8192 P7 cores clusters
Pre/Post-Processing clusters (PPP)	Two 640 X86 cores custom clusters
Global Parallel Storage (Site-Store)	CNFS and ESS clusters
Near-Line Storage (HP-NLS)	StorNext based archiving cluster
Home directories	Netapp home directories

As well as

- Hosting of the Solution
- High Performance Interconnects
- Software & tools
- Maintenance & Support
- Training & Conversion support
- On-going High Availability

ECCC Supercomputing Procurement Requirements

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

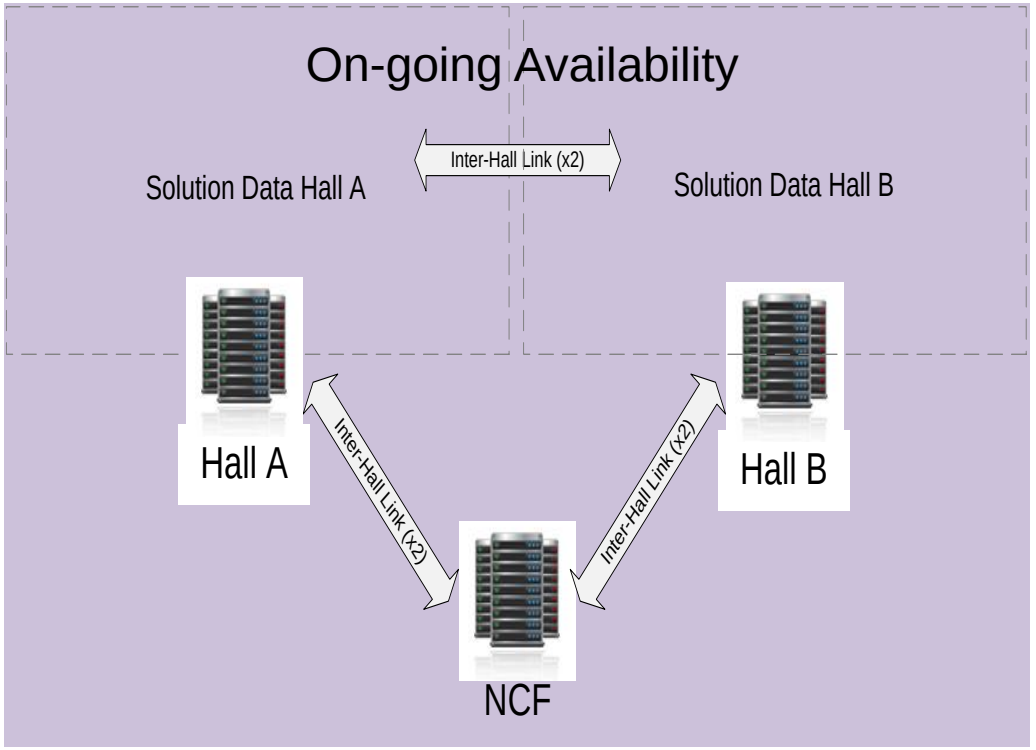

Outcome

- IBM was awarded the contract
  - Evaluation based on benchmark performance (Forecast Day's per Hour) on a fixed budget
- Proposal for initial system
  - Supercomputer: Cray XC-40, Intel Broadwell, CLE 5.2, Sonexion Lustre Storage
  - PPP: Cray CS-400, Intel Broadwell, RHEL 7*
  - Site-Store and Homes: IBM Elastic Storage Server (ESS, GPFS-based)
  - HP-NLS: based on IBM High Performance Storage System (HPSS)

Sizing

- Computing
  - About 37,000 Intel Broadwell cores per Data Hall
    - Super and PPP combined
- More than 40PB of disk storage
  - 3.6 PB scratch storage per supercomputer (one per data hall)
  - 18 PB site store per data hall
  - 1.1 PB disk cache to the archive per data hall
- More than 230 petabytes of tape storage (two copies)

The Newest Addition to a Long History

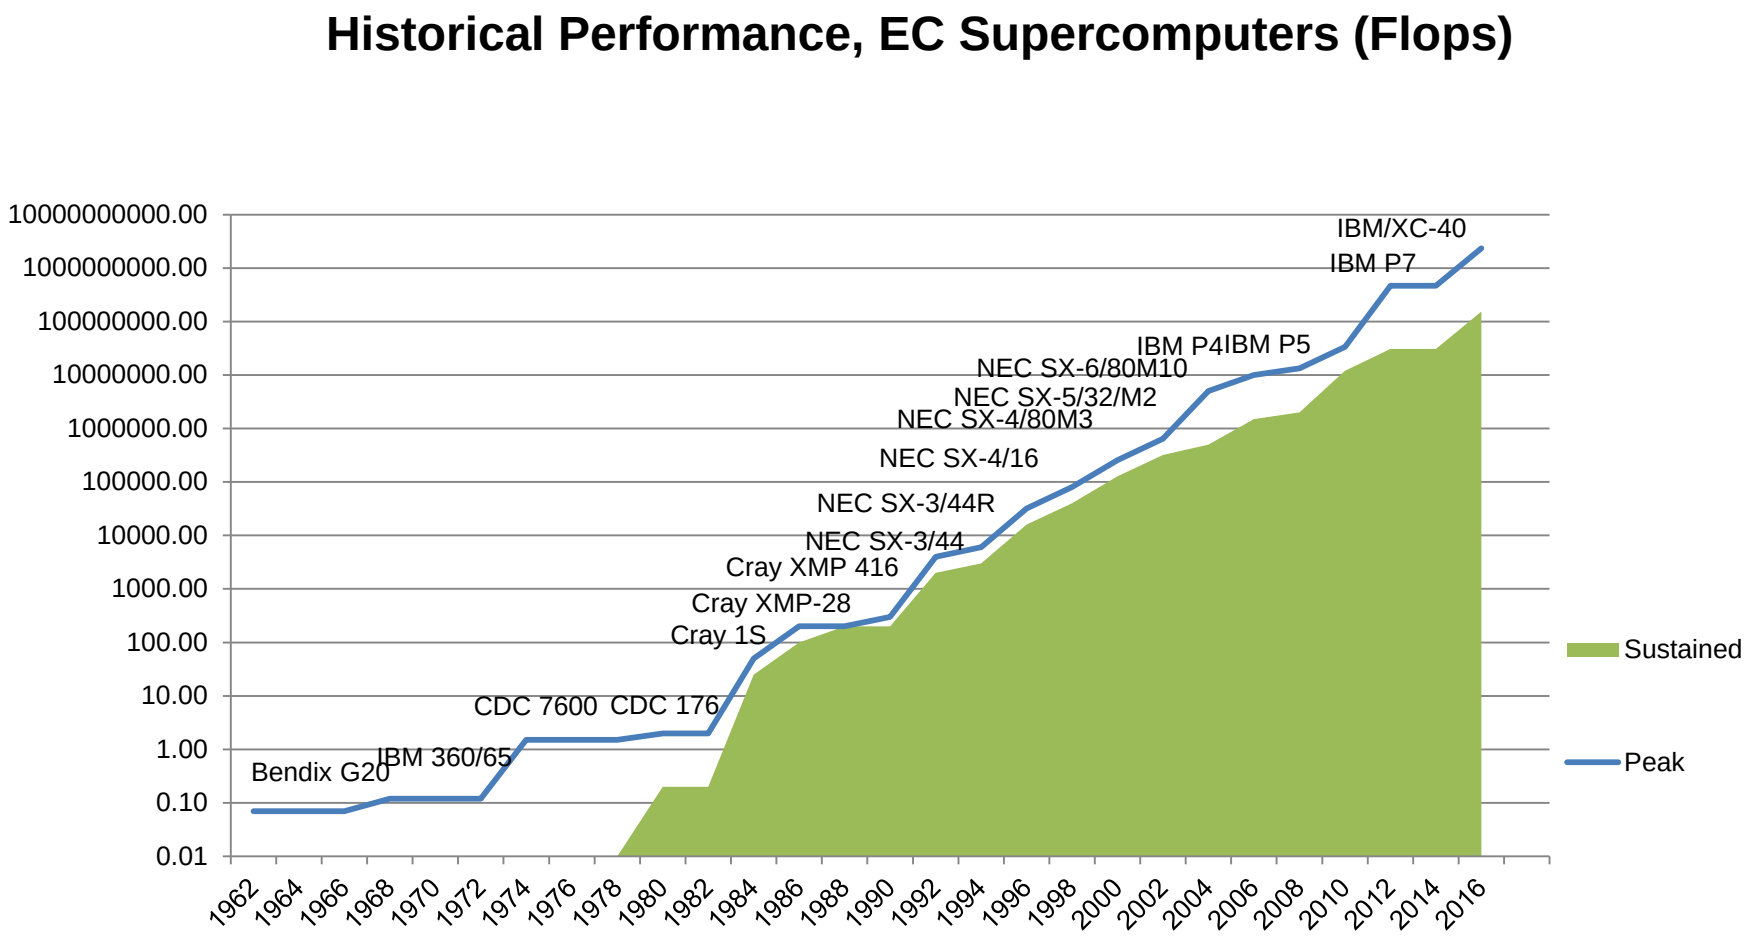

Resulting Architecture

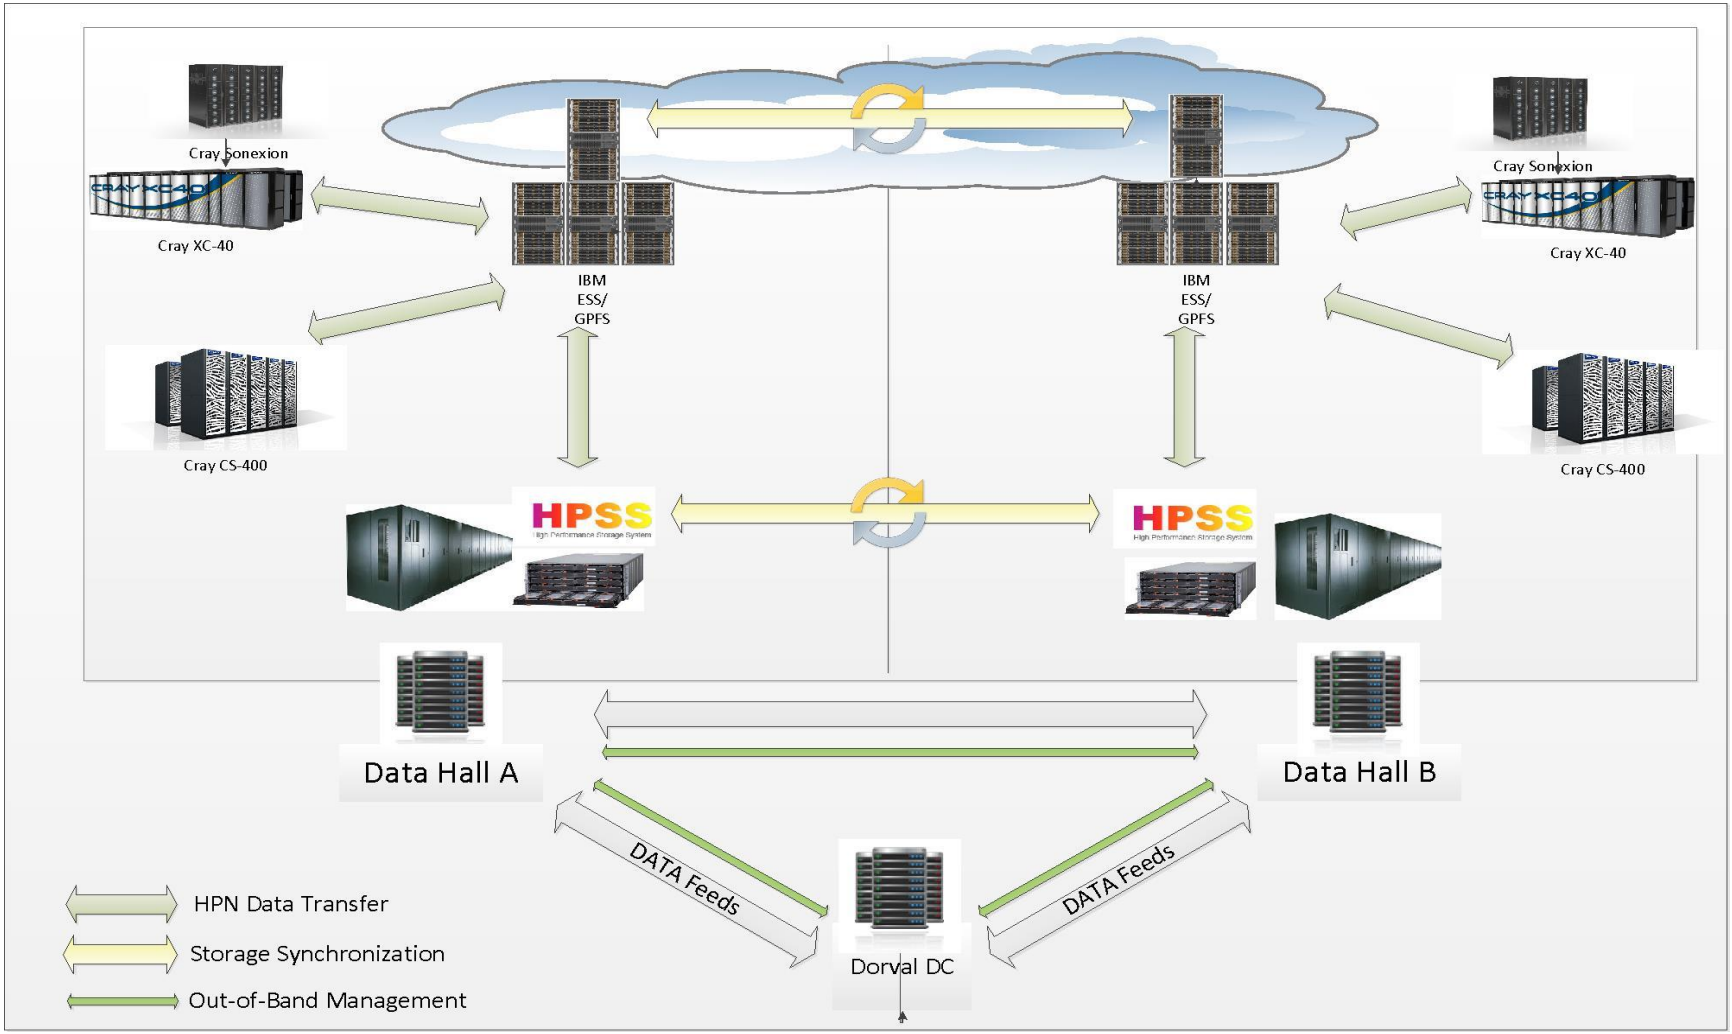

Challenge and Status

- Change the Supercomputer clusters, PPP clusters, archiving system and homes. All at once. Never been done for us.
  - A lot of preparation work was done ahead of time
    - Most codes were already been ported to Intel architecture
    - Our General Purpose Science Clusters available for PPP migration work
      - Linux containers are being leveraged to smooth the transition
  - We're almost there
    - Final cycles done
    - Parallel runs starting
    - Just a few steps left

Q&A

Thank you.