Raytheon

*Customer Success Is Our Mission*


# Adopting HTCondor at Raytheon

Michael V. Pelletier
Principal Engineer
Information Technology
May 2015

# About Michael Pelletier

Raytheon

- Raytheon IDS IT Organization: March 2, 2009
  - *Program Execution & Business Work Environments*
  - *University of Wisconsin Flexible Option 2014-*
- Nortel Networks: 2000-2008
  - *via Alteon Web Systems*
- Taos: 1998-2000
  - *Sysadmin consulting for Stratus, others*
- TechTeam Global: 1992-1998
  - *Ford Motor Company consulting*
  - *Established WordPerfect's first contract helpdesk*
- University of Michigan Engineering: 1988-1992
  - *Sysadmin for CAEN tech support organization*
- US Naval Academy Summer Seminar: 1986
- M-net, a UNIX-based BBS: 1986-1996
- Commodore PET, VIC-20, and 64: 1982-1990
  - *BASIC and 6502 assembly language*

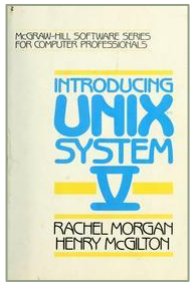


**1986**

**1982**

# About Raytheon

**The Missile Defense Center**

*Photo by Cranshaw Construction in Woburn, Massachusetts*

**Rebecca Rhoads**

*CIO, Executive Diversity Champion*

*#6 ExecRank Top 50 Female CIOs*

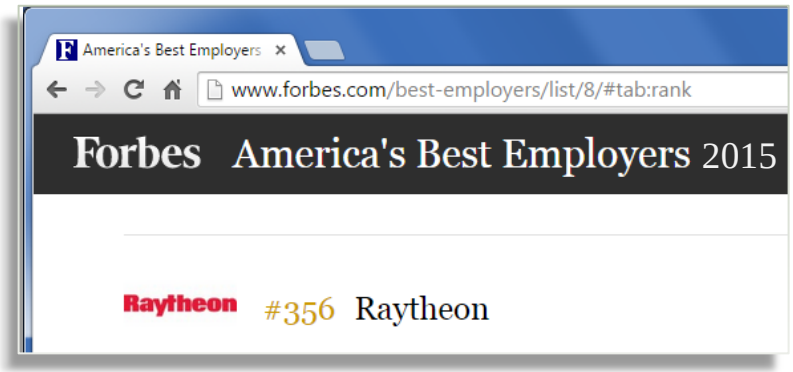


CR's 100 Best Corporate Citizens 2015

HUMAN RIGHTS CAMPAIGN

TOP GREEN IT ORGANIZATIONS

COMPUTERWORLD
BEST PLACES TO WORK in IT
2013 Employer Profile
NO. 58
Raytheon
Waltham, Mass. [New England]

**www.rayjobs.com**

# About Raytheon

- **Raytheon Company**
  - Founded in 1922
  - $23 billion in 2014 sales
  - 61,000 employees worldwide

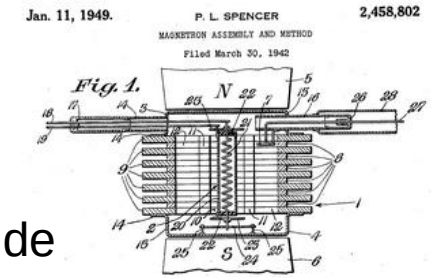


- **Integrated Defense Systems**
  - Headquartered in Tewksbury, Massachusetts
  - Broad portfolio of weapons, sensors, and integration systems across multiple mission areas including **air and missile defense radars**; early warning radars; naval ship operating systems; command, control, communications; air traffic systems


- **Information Technology**
  - "Today's extensive High Performance Computing needs require us to continuously push the envelope to create leading-edge technical processing solutions for complex systems and program challenges."

**Nearing a century of technology and innovation**

# Raytheon's Vision

## VISION

One **global team** creating **trusted, innovative solutions** to make the **world** a **safer** place.

# Making the World A Safer Place


**Most above are or contain Raytheon equipment**

Raytheon

# A Safer Place in Ascent Phase


**Forward-based AN/TPY-2 radar**

Raytheon

# A Safer Place in Midcourse Phase

**"Should one get airborne and come at us, I'm confident that we'll be able to knock it down."**
**Adm. Bill Gortney, Cmdr NORAD**
**April 7, 2015**

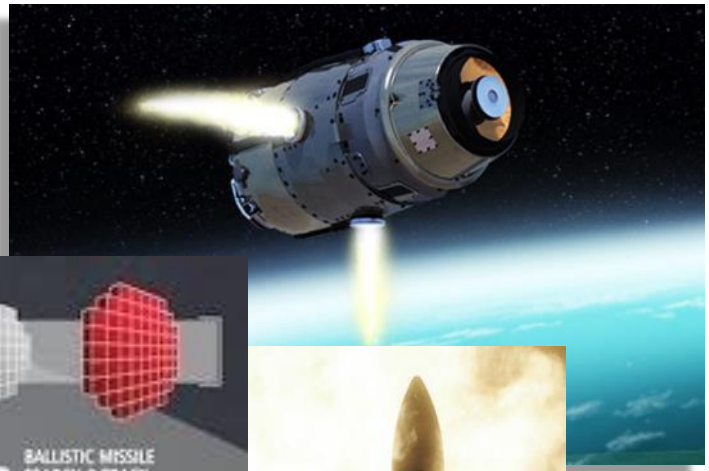

***SM-3 Kinetic Interceptor (Hit-to-Kill)***

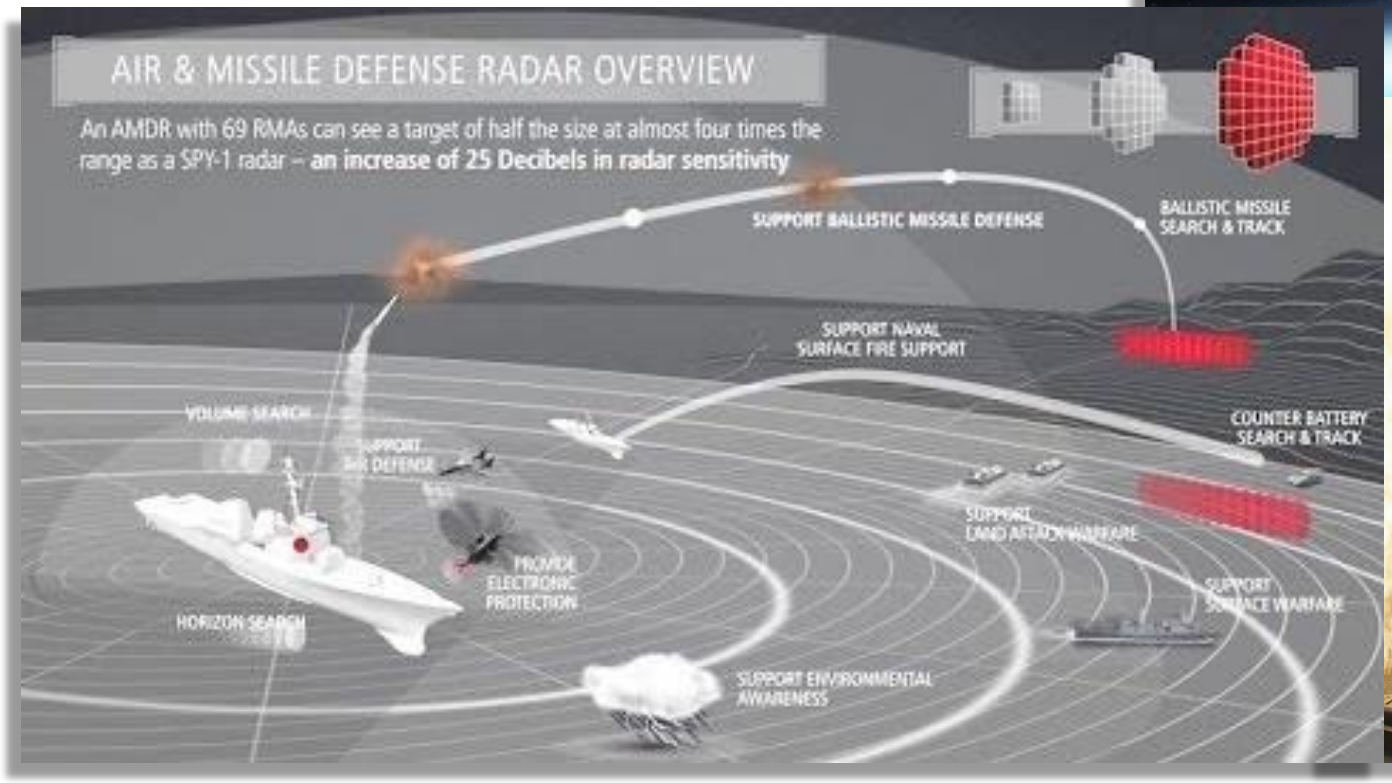


**AMDR will replace Lockheed's AN/SPY-1 in Aegis**

# A Safer Place in Terminal Phase


Threat missile before…

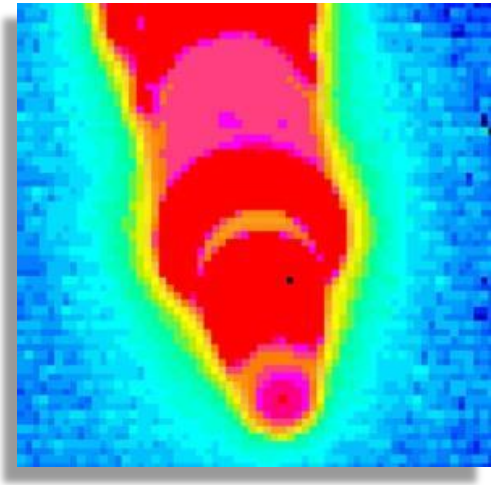

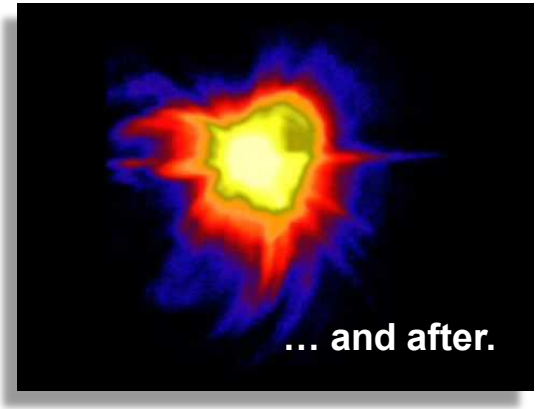


**THAAD with AN/TPY-2 in terminal mode**
**(Terminal High-Altitude Area Defense)**

# A Safer Place at Close Range

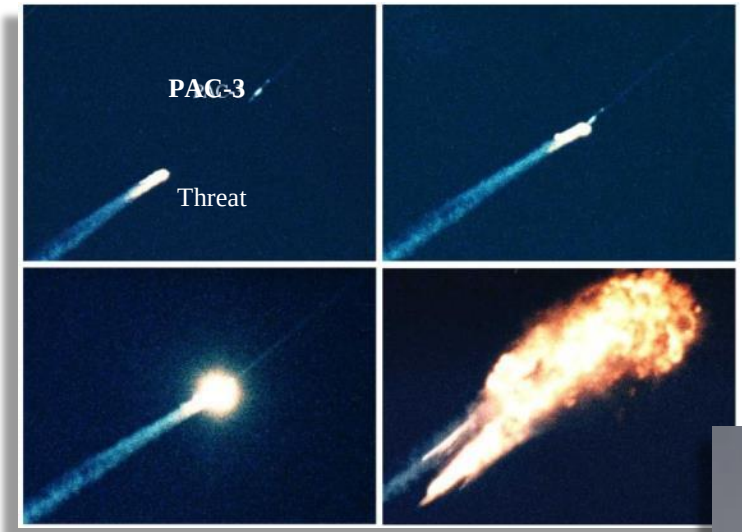


*Patriot PAC-3*
*flight test intercept*
*October 14, 2000*

RT

**Turkey finds sarin gas in homes of suspected Syrian Islamists – reports**
*May 30, 2013*

Patriot defending a city near the Syrian border during the Syrian Civil War
February 2013

**Patriot Systems Improvement Program**

# Common Thread: Modeling & Simulation


*Raytheon has been deeply involved with **all phases of modeling and simulation**, from the design and development of models and simulation systems to their use in real-world, mission-critical applications.*

*…**simulations** play a central role in optimizing the value of what are **typically expensive field tests**, and in some cases simulations actually reduce the amount of field testing needed.*

**Mark Russell**
*Vice President of Engineering, Technology & Mission Assurance*

**Modeling and simulation is a matter of life and death**

Raytheon

# Excellence Saves Lives


**Evolving threats demand**
**faster, better, *more* simulation runs**

# Simulation Workflow

Trajectory Generator, Xpatch, OMP Runs

CX-EDT, ETESIM, RAMS Runs

Input Generation

Simulation Processing

Output Analysis

Reduction, Forming, Plotting, Analysis Runs

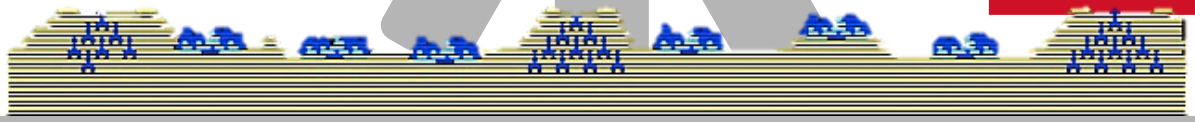

**Opportunities for improvement in all areas**

# Foundation: AN/TPY-2 Program 2010 to 2012

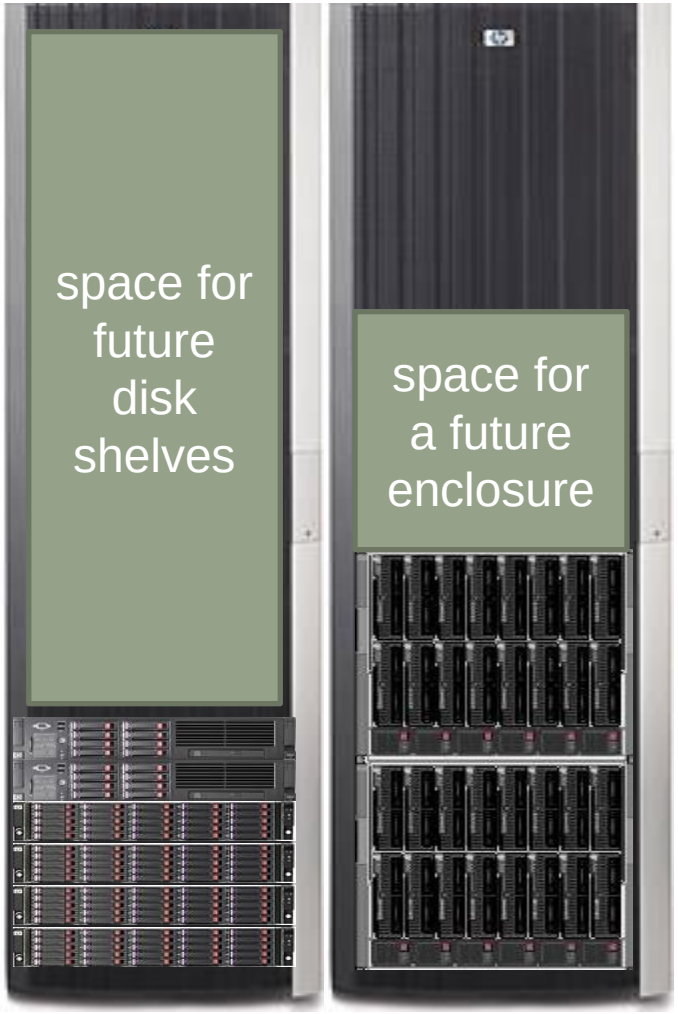


- **Open-source HTC alternative**
- **Memory was not configured as a "consumable resource"**
- **No usage shares configured**
  - Queues were first-come, first-serve
- **Many hard dependencies**
  - Menagerie of engineering workflow tools needed specific paths and different qsub capabilities.
- **Poor network performance**
  - Tuning, hardware compatibility
- **Lots of non-scalable jobs**

**The good news: infrastructure with room to grow**

# The HTCondor Egg

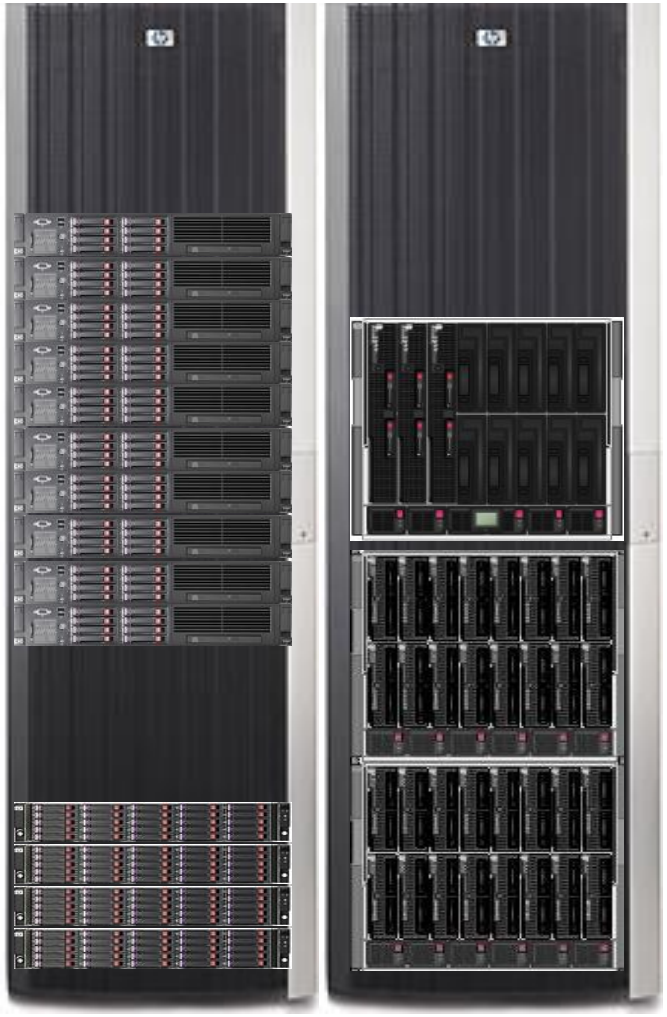

- **Non-scalable code led to memory oversubscription**
  - Existing execute node memory was doubled to 96GB
  - Added 3 BL620c 40-HyperThread 512GB blades in mid-2012
- **First HTCondor in Early 2013**
  - Spare DL380 as Central Manager
    - HP Insight CMU / Diskless boot with local-disk scratch space
  - 3 new 512G BladeSystems
  - 7 spare DL380 ProLiant servers
  - Red Hat Enterprise Linux 6

**Initial pool had 288 cores, ~35% of original system**

# Struggling to Switch Platforms

- **Migration Obstacles**
  - MATLAB DistComp Toolbox used one qsub per task
  - Homebrewed sge_helper.py used qsub –t start:step:end syntax
  - Embedded dependencies and assumptions inside layers of code
  - Resistance to porting code for unproven platform

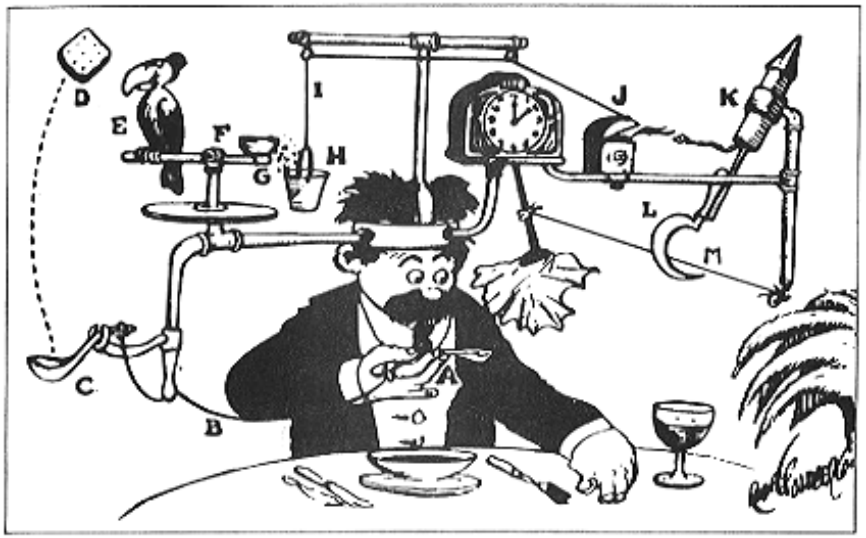


- **Wrote a Perl script from qsub man page over a weekend**
  - The January 2013 condor_qsub didn't fit all our use cases
  - Command-line compatible Perl script worked for everything
  - Saved the migration from the code-porting impasse
- **CXSim X-band Simulation**
  - First to migrate to the new pool

**Untangling legacy code is an ongoing challenge!**

# Immediate Benefits with No Code Change

- **Improved Job Management**
  - HTCondor's ClassAd matching model and priority system eliminated first-come, first-serve queueing and memory oversubscription.
- **Higher Throughput on Same Hardware**
  - HTCondor management gave fewer failed runs and less swapping
  - Big improvements in kernel performance between RHEL5 and RHEL6
    - 10-15% better job runtimes even with no other differences
  - Properly-tuned RHEL6 kernel and network multiplied system performance
- **Resource Utilization Clarity**
  - RHEL6 Control Groups and OOM killer help prevent system hangs and crashes, allowing identification of resource problems on a live system
  - HTCondor job logs provided easy analysis of resource usage patterns
- **CollectL and CondorView Clarity**
  - Better understanding of job and system behavior, and failure modes

**HTCondor cleared the brambles, exposed the pitfalls**

# Leveraging HTCondor Capabilities

- **Monte Carlo CXSim / CX-EDT Submission Script**
  - First job-submission tool which I converted to HTCondor
  - Huge increase in queueing rate to go with run speed increase

One qsub per task.
Startup delay to interleave users.
Separate script for run area setup.
NFS-based run area.
Unmanaged output delivery.

queue 1000
Seed list indexed with $(Process)
Run area setup by exec node
$_CONDOR_SCRATCH_DIR
transfer_output_files

**HTCondor out-produced old pool with 1/3 the CPU**

# Hatching the Egg, Growing the Flock

- **Converting existing compute nodes to HTCondor**
  - With less CXSim demand, more old systems idle
  - HP Insight Cluster Management Utility diskless-boot made migration easy
- **Network tuning**
  - Wrote interface bonding into system-config-netboot
  - Improvements to BladeSystem configuration, firmware
  - Tuned TCP stack and found a 1.25Gb bottleneck in the TuboIron 24X 10Gb switch (still a mystery)
- **Converting more tools with qsub**
  - Having a drop-in replacement simplified migration
  - Old MATLAB submit function needed help
    - Had to set LD_LIBRARY_PATH when calling HTCondor binaries in qsub
    - Didn't get around to writing a HTCondor-native MATLAB interface until early 2015

**Trajectory Generator tool was the next to migrate**

# Trajectory Generator

- Government-supplied black-box software can't be modified
- "Laydown file" contains hundreds of launch and aim points which can be processed independently

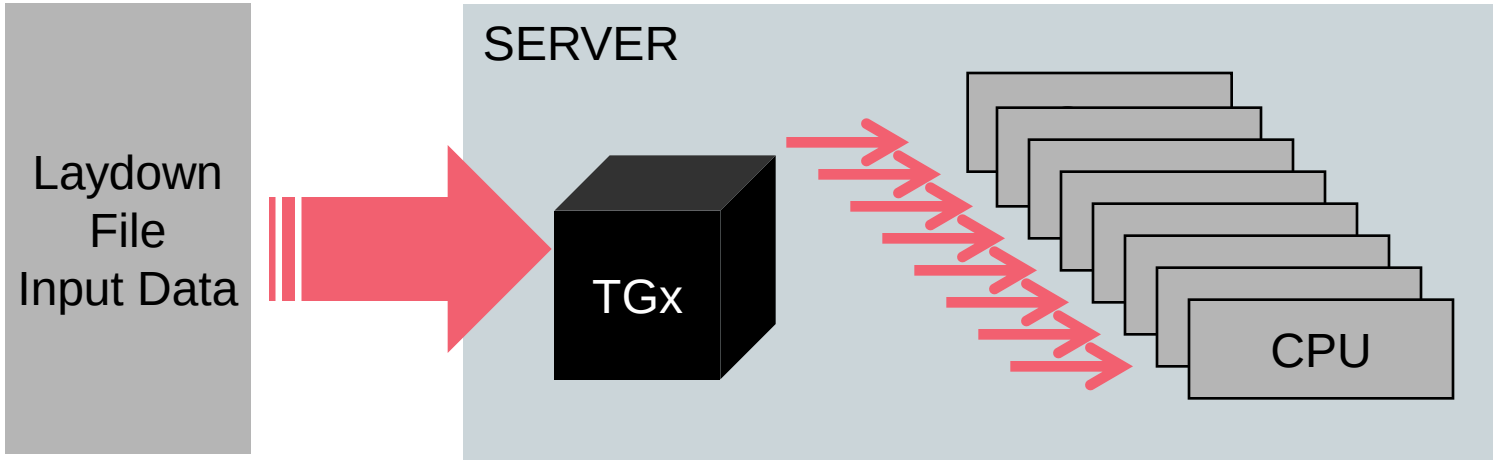


- TGx by default used all cores in the machine, crowding out other processes as it worked through the list of trajectories
- 320 trajectories on a 24-core server required 13 hours

**Needed to Turn 1 320-Set Run Into 320 1-Set Runs**

# Trajectory Generator Automation

- Wrapper script splits laydown file into one-trajectory chunks:

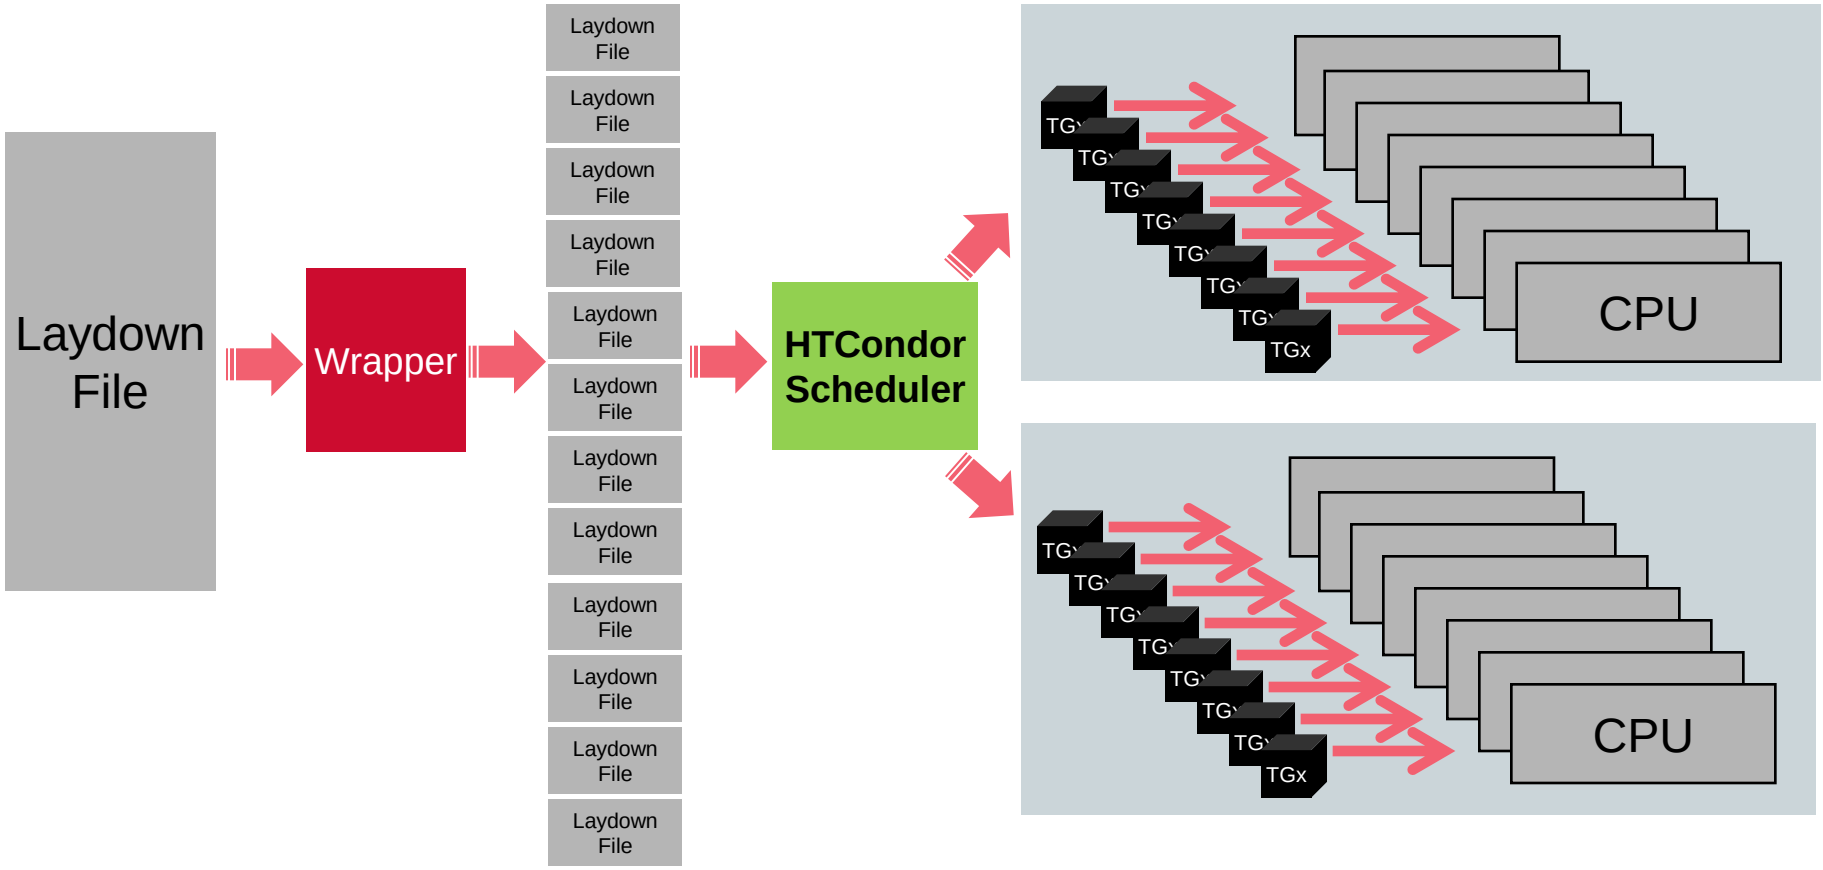


HTCondor's clear architecture made implementation straightforward.

**Can run all 320 trajectories in *45 minutes!***

# Old Platform vs. New HTCondor

AN/TPY-2 Compute Nodes in 2012


AN/TPY-2 HTCondor in late 2013

**Clarity and logical consistency delivered results**

# Living Raytheon Values & Mission

- **Collaboration**
  - Sharing knowledge enabled optimal solutions
- **Innovation**
  - Challenging the status quo of inefficient code
  - Devising better ways to run critical workflows
- **Accountability**
  - Anticipating the needs of IT's engineering partners for computing capacity
  - Doing what it takes to deliver the promised results to Raytheon's customers
- **Information Technology Mission**
  - Fuse business insight, innovative application of information technology and collaborative relationships for program and partner success.


**IT / Engineering partnership was key to success**

# AN/TPY-2 Program's HTCondor Success

Recognized with the
**2013 Excellence in Engineering & Technology Award**
(Raytheon's Highest Technical Honor)

Awards dinner was held at the National Air & Space Museum in Washington, D.C.

**Key factor in award was IT / Engineering collaboration**

# Air and Missile Defense Radar & RAMS

- **New program began October 2013**
  - HTCondor implementation started January 2014
- **R**aytheon **A**ir and **M**issile **S**imulation
  - Java with CUDA offloads for key calculations
  - Runs used to be managed by home-brewed scripts on desktop Linux machines
  - HTCondor saved them from reinventing LSF
- **HP WS460c Blades with K4000 GPUs**
  - Dense: 8 servers and 16 GPUs in 10U enclosure, and can get even more dense with up to 48 K3100M GPUs
  - HTCondor GPU Advertising
    - Used 8.1dev condor_gpu_discovery in 8.0.6, before upgrade to 8.2.6
    - Adapted ClassAds to reflect RAMS requirements

NVIDIA®

**Multiple time-consuming workflow steps eliminated**

# AMDR HTCondor Results

"In support of ongoing analysis for the Performance Assessment Working Group and the future AMDR Systems Critical Design Review, we are now able **to track our progress towards meeting design goals** by examining our build performance against a large subset of our required performance space **in just a few hours**. This allows us to do a thorough regression test on a nightly basis to **catch bugs as soon as they are introduced.**

"Additionally, we are now able **to run our full analysis space of thousands of cases in less than a week,** allowing time for design iteration. The **ease** of using the **HTCondor** system with our existing tools has facilitated our design and analysis work and we are heavily dependent on it for our work going forward." (November 2014)

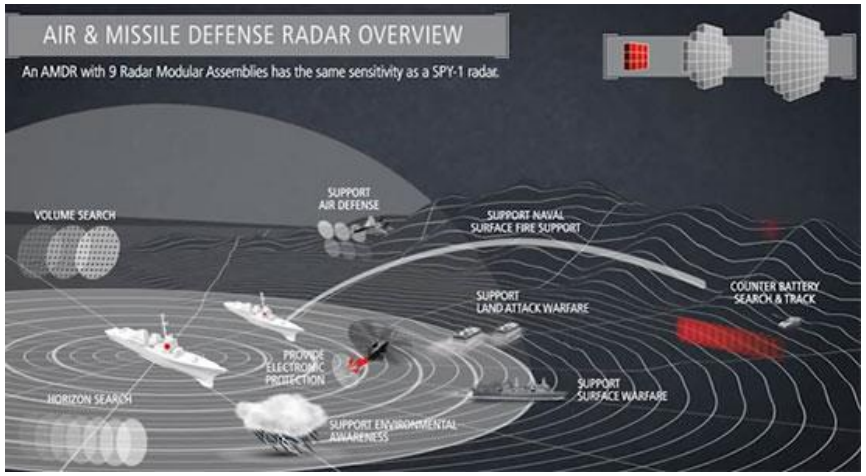


**– An AMDR Lead Developer, Principal Engineer, and fellow EiET Awardee**

**Enabling Agile excellence for AMDR**

# Breaking News: AMDR CDR Success!

**TEWKSBURY, Mass., May 12, 2015 /PRNewswire/**
The U.S. Navy and Raytheon Company (NYSE: RTN) have completed the AN/SPY-6(V) Air and Missile Defense Radar (AMDR) critical design review.

The outcome confirms Raytheon's design and technologies as mature, producible and low risk; on track to meet all radar performance requirements, on schedule and within cost.

**HTCondor delivers, for Raytheon and the nation!**

# Patriot Systems Improvement

- **HTCondor for Monté Carlo Simulation Runs**
  - "HTCondor was **far simpler to install and configure** than the system which it replaced. It was also very interesting to learn that we could install it on our workstations as well, and add them as opportunistic compute nodes for our cluster.

    My users are **very happy** with it. The transition to HT Condor was much easier for them then I expected. Michael and I had them up and running in no time!"
    (November 2014)

    **– Patriot Systems Improvement Lead Systems Administrator, IDS Information Technology**

**HTCondor scalability is a key discriminator**

# MATLAB Distributed Compute Engine

- **Used qsub in MATLAB generic scheduler at first**
  - HTCondor concurrency limits invoked via qsub "-l" command line option which had to be written into the qsub wrapper
  - Sluggish submissions – about 30-40 per minute
- **Wrote HTCondor-native interface in February 2015**
  - Maps job.numberOfTasks from MATLAB scheduler object to "queue" in submit description, so it can submit thousands of jobs in a few seconds instead of a few hours as we had with one qsub per task
  - Kudos to the MathWorks Parallel & Distributed Computing product team for their assistance in developing the new interface! Work is still in progress…


**Accelerating the pace of Job/Task submissions**

Raytheon

# Ongoing Improvement

AN/TPY-2 Compute Nodes in late 2013



AN/TPY-2 HTCondor in 2015

**Still some construction zones, but excellent progress**

# Future Potential

- Patriot Systems Improvement
  - Investigating HTCondor on the Windows platform
  - Windows / Linux heterogeneous pools
- Raytheon Missile Systems in Huntsville
  - Replace existing software for multi-spectral scene generation and other jobs
- Upgraded Early Warning Radar
  - Replacement of existing software and Red Hat 5 and 4
- Sea-Based X-Band Radar
  - Replacement of existing software and Red Hat 5
  - Re-architecture of existing hardware in light of AN/TPY-2 lessons learned
- Unclassified projects and programs
  - Investigation of opportunistic and VM execute nodes on infrastructure systems
  - Internal & external cloud services for proposal and program execution

**Interest from many different parts of the company**

# The Right of Self-Defense

***The right of self-defense never ceases. It is among the most sacred, and alike necessary to nations and to individuals…***

**President James Monroe**
November 16, 1818

**Fifth President of the United States**

# Making the World a Safer Place

***I call upon the scientific community in our country, those who gave us nuclear weapons, to turn their great talents now to the cause of mankind and world peace, to give us the means of rendering these nuclear weapons impotent and obsolete.***

**President Ronald Reagan**
March 29, 1996

**Fortieth President of the United States**

# Questions and Comments