



# Machine Learning Techniques for Improving Flash Endurance

Conor Ryan & Joe Sullivan  
CTO - Software/Hardware

[Conor.Ryan@NVMdurance.com](mailto:Conor.Ryan@NVMdurance.com) & [Joe.Sullivan@NVMdurance.com](mailto:Joe.Sullivan@NVMdurance.com)

# Take Home Messages



- ▶ 3D flash is too complex to trim effectively with current methods
- ▶ NVMdurance Machine Learning scales to meet the challenge
  - ▶ Marriage of simulation and real world testing
- ▶ Fully automated trimming used on two drives at FMS
  - ▶ NVXL (stand no. 801)
  - ▶ Altera-Intel/MobiVeil (stand nos. 120 and 610)
- ▶ Full toolkit and reference design available for SSD makers
  - ▶ See us at stand 829

# Take Home Messages




- ▶ 3D flash is too complex to trim effectively
  - ▶ NVMdurance Machine Learning scales to:
    - ▶ Marriage of simulation and real world test
  - ▶ Fully automated trimming used on two:
    - ▶ NVXL (stand no. 801)
    - ▶ Altera-Intel/MobiVeil (stand nos. 120 and 121)
  - ▶ Full toolkit and reference design available for SSD makers:
    - ▶ See us at stand 829
- ▶ Results

  - ▶ 3-10X increase in endurance
  - ▶ Application-specific trimming
  - ▶ Running in drives right now

# Flash Trimming



- ▶ The art of finding flash parameters
  - ▶ To achieve reasonable specification for broad appeal
  - ▶ To specific/extreme requirements
- ▶ Many parameters interact with each other
  - ▶ Satisfy one criterion (e.g. low BER)...
  - ▶ Violate another (high tProg and tErase)




# It just got harder



- ▶ 3D NAND has an order of magnitude more complexity
- ▶ Machine Learning can model and **automatically trim flash**
  - ▶ Flash can be trimmed for different applications
- ▶ Flash vendors don't **optimize** flash, they make *it good enough for broad markets*
  - ▶ *Achieve X cycles with 3/12 months retention*





# Two Pronged Approach



- ▶ NVMdurance Pathfinder
  - ▶ Discover parameter sets to satisfy goals
  - ▶ Discover multiple sets of parameters, each tuned for a particular time of life for the Flash
- ▶ NVMdurance Navigator
  - ▶ Lightweight software that runs on the SSD controller
  - ▶ Exploits Pathfinder-derived parameters and deals with variability
  - ▶ Does so by changing LUN parameters based on *health* indicators (RBER/thresholds/timing/history)
- ▶ Best results are found when both are used; however, either can be used on its own

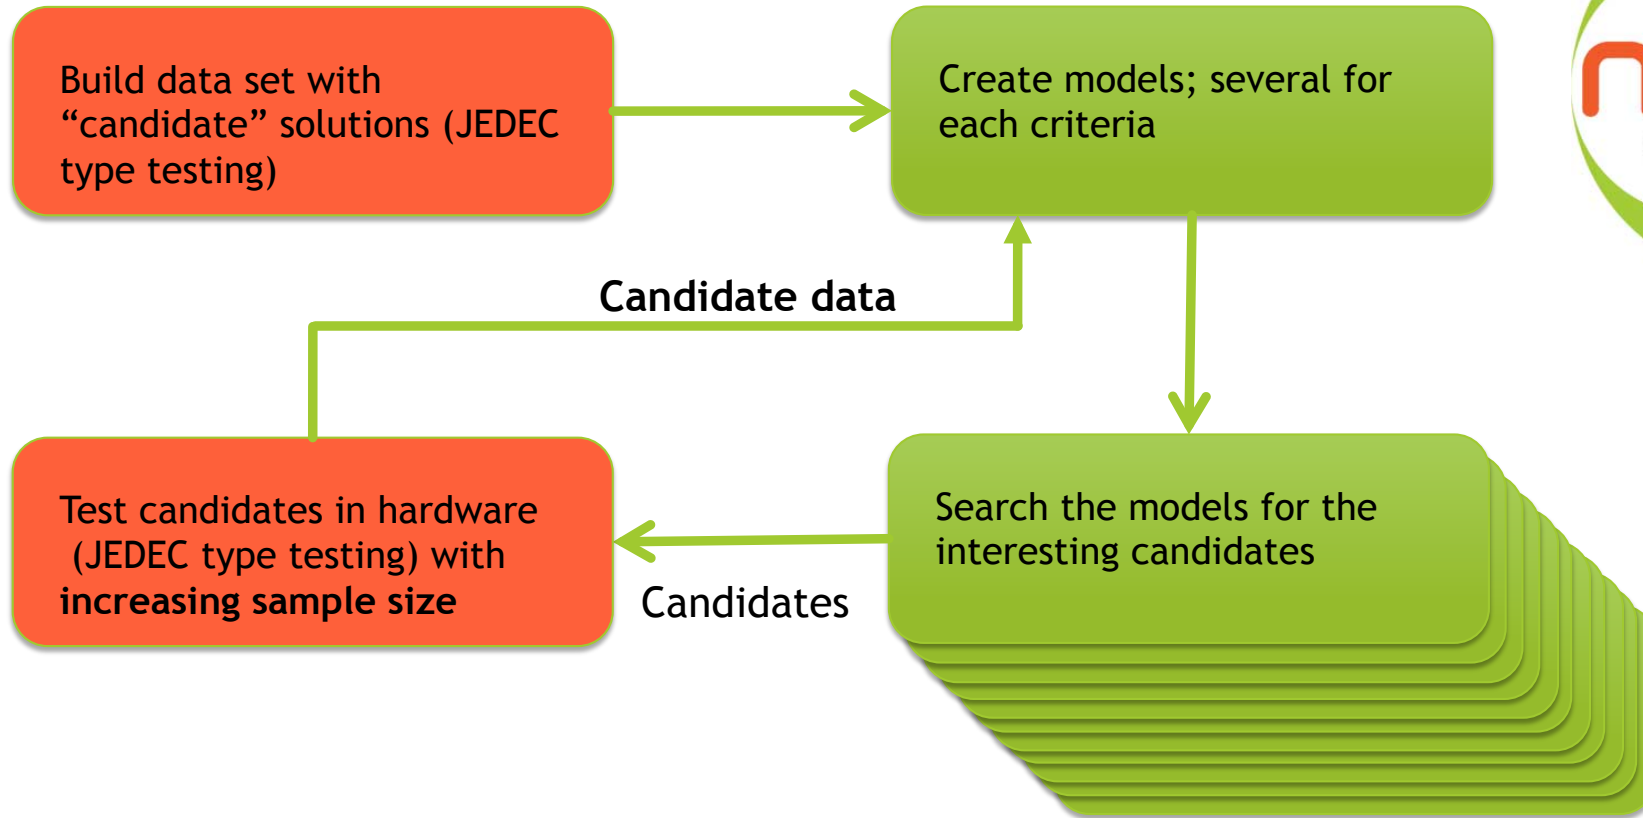
# Machine Learning - NVMDurance Style



- ▶ Machine Learning discovers patterns in big and noisy data
  - ▶ Stores knowledge that is
    - Searchable
    - Incremental
  - ▶ We're learning how parameter sets perform on test criteria
- ▶ Search
  - ▶ Find best parameter set using the models as *surrogate testers*, given
    - Noisy data and possibly inaccurate results
- ▶ Validation
  - ▶ Test the parameter sets in real hardware

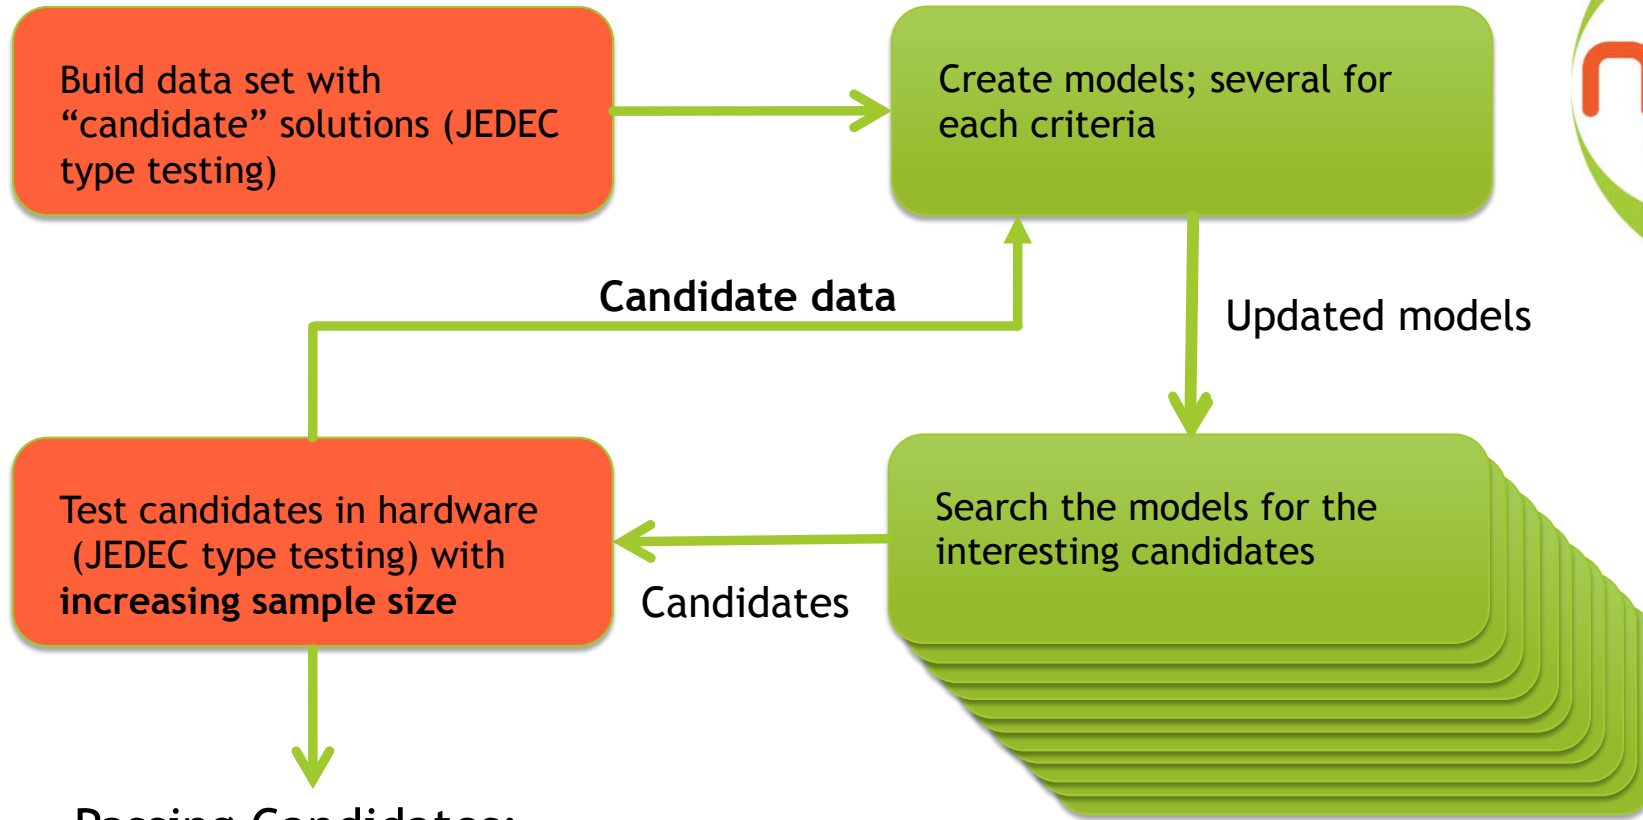


# Data Flow



200 8-bit registers =  $2.56 \times 10^{18}$  candidates

# Data Flow



200 8-bit registers =  $2.56 \times 10^{18}$  candidates

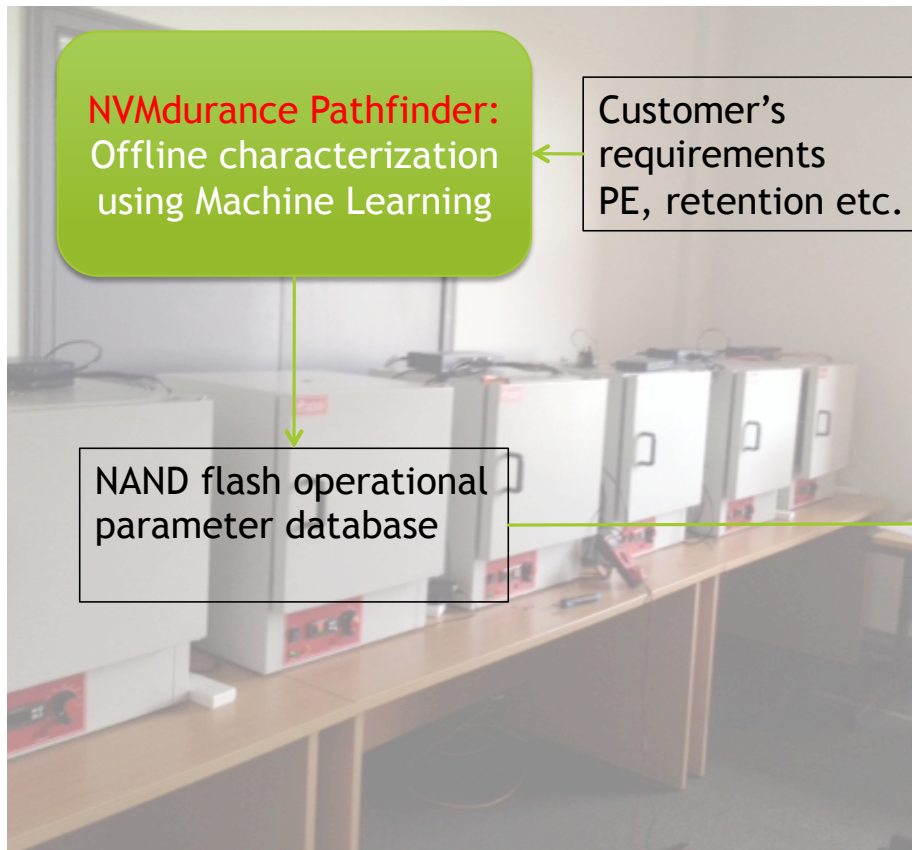
# Scaling

- ▶ Scaling factor from hardware tests to software search is at least six orders of magnitude
  - ▶ 20 hardware tests can lead to 20 million virtual tests
- ▶ But...
  - ▶ Simulation is cheap and fast; this is already increasing
  - ▶ “Force multiplier”: simulation dramatically improves the power of Machine Learning
  - ▶ Hardware validation enforces sanity checks

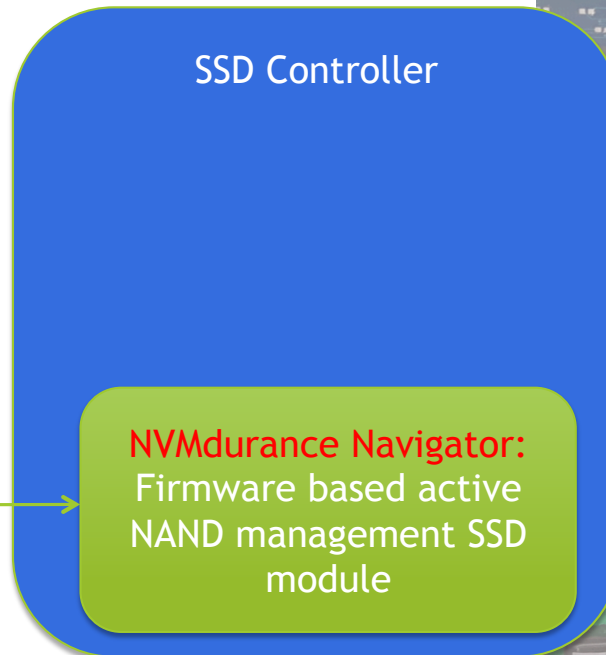


# NVMdurance Patented Process

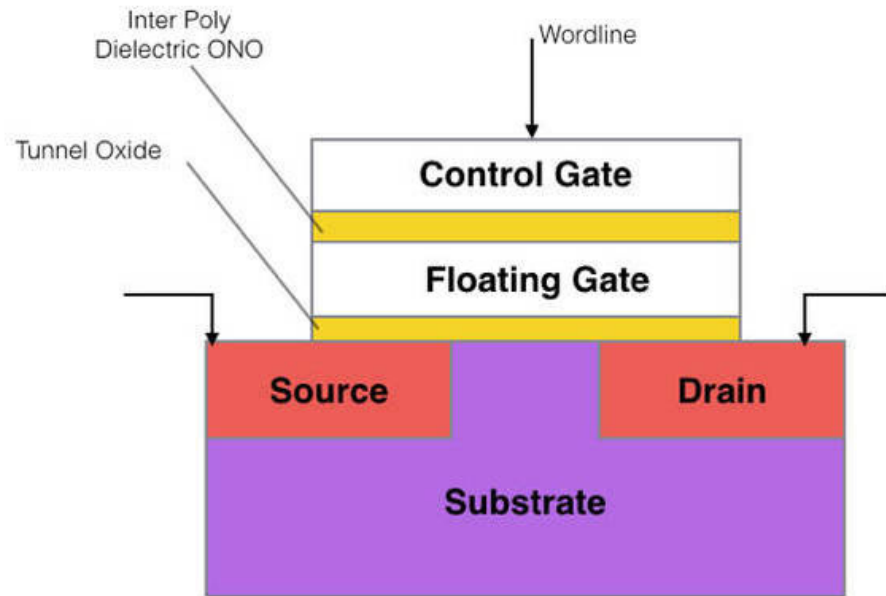
Offline | On Controller



Flash Memory Summit 2016  
Santa Clara, CA



# Flash wear-out mechanics



- Large voltages used to push electrons on and off floating gate
- Electrons passing through tunnel oxide damage it, so are more likely to drift off the floating gate
- Electrons get stuck in tunnel oxide; obstruction causes erase difficulties

## How and Why does it work

- ▶ Off line characterization discovers optimal operational parameters for each of up to 5 life stages for specific retention periods
- ▶ **NVMdurance**: Each parameter set reduces wear by applying only the charge required to each storage element, to make the retention figure desired by the application at the PE for the end of that stage
- ▶ **The NAND FAB**: The factory parameters applies charge (throughout life without change) required to make the Jedec retention figure at the end of PE



# Example: MLC 1 years retention 5k PE cycles



## In the FAB Solution

For every PE cycle from 0 to 5k  
We must always pass enough charge  
such that at 5k PE the cells will have  
bit flips < ECC rate after 1 years retention



**In NVMdurance Solution**  
For PE cycles from 0 to 1k  
Pass on enough charge such  
That at 1k PE the cells will  
have bit flips < ECC rate  
after 1 years retention



## In NVMdurance Solution

For PE cycles from 1k to 2k  
Pass on enough charge such  
that at 2k PE the cells will  
have bit flips < ECC rate  
after 1 years retention Etc.



## Why use this approach?....

- ▶ NAND media last at least 3 times longer when powered by NVMdurance
- ▶ Number of LEs required lower by reduced ECC needs
- ▶ LDPC Hard decode (or BCH) give a predictable, tail latency free response times
  - ▶ No need for soft LDPC



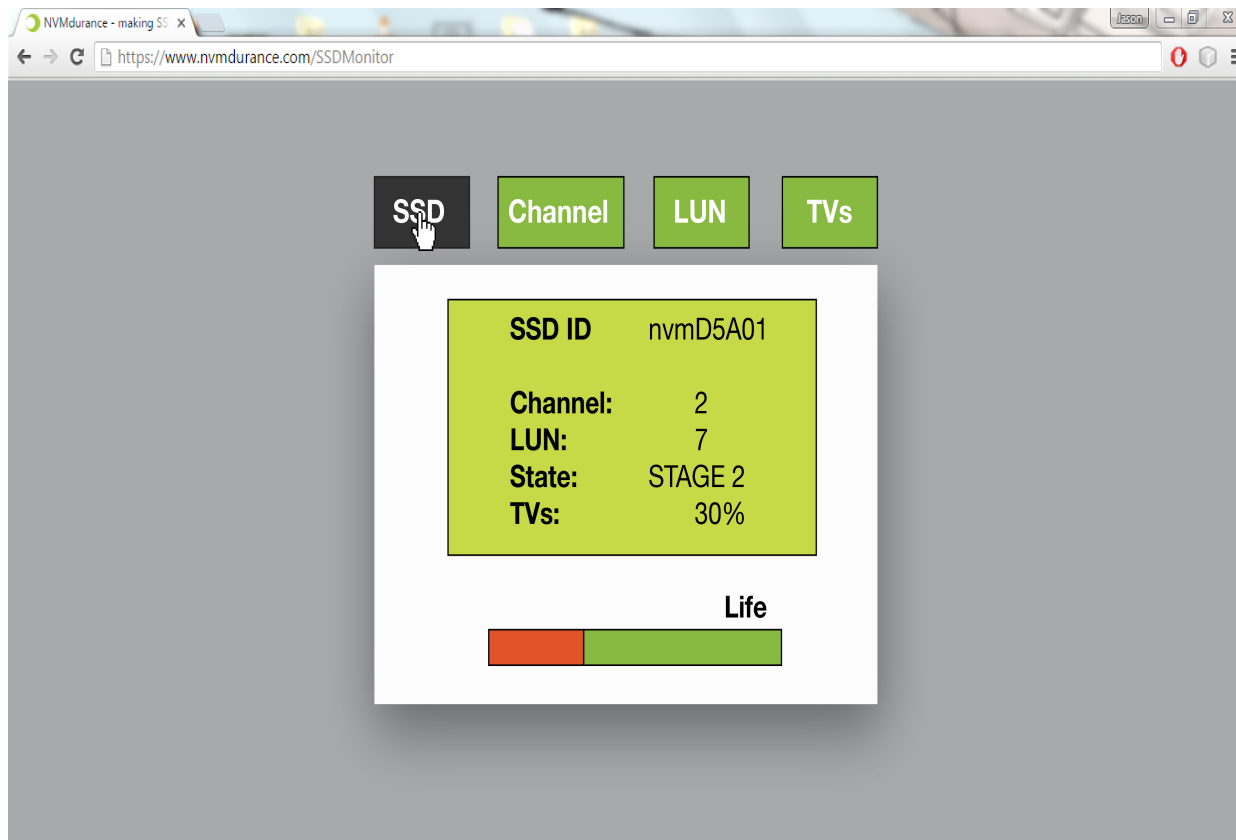


## Why use this approach?....

- ▶ Each SSD is highly configurable in the field and may be deployed or redeployed in any number of ways
  - ▶ e.g. From ‘Read Intensive Zero Tail Latency’ to ‘Archive, Long Retention’ or anything in between
- ▶ Comprehensive reporting of life stages and remaining life estimates
- ▶ Simple upgrade path for new devices or as firmware or FPGA-ware improves.
  - ▶ a simple database swap



# SSD Real-Time Extensive Life Reporting



- SSD life may be monitored by SSD, per Channel or per LUN
- SSD may be re-tasked by swapping of LUN operational parameters provided by NVMdurance

## What we are showing today at FMS


- ▶ NVMdurance Alaric Development board SSD POC reference design
  - ▶ NVMe over PCIe
  - ▶ 4 channels, single LUN per channel, 1 Gbyte total
  - ▶ 40 bit BCH ECC
  - ▶ NVMdurance Navigator active flash management (life extension 5X)
  - ▶ NVMdurance operational parameters database
  - ▶ Planar TLC devices
  - ▶ NVMdurance Navigator is demonstrable on separate NAND test head



# NVMdurance Alaric Dev. board SSD POC

4 channels, single LUN per channel, removable media

Altera Arria 10 running 40 bit BCH ECC, channel controllers NVMdurance Active Flash Management



TLC NAND

NVMe over PCIe

# The NVMdurance Advantage



- ▶ The operational parameter are tuned to your application and not the vendors highest sales pipeline
- ▶ NVMdurance Navigator manages the parameters, the optimal read poles, and adjusts for wear and NAND production variation
- ▶ Retuning SSD in the field is a simple matter of switching parameter database values (in planar MLC this is about 60 bytes)

# NVMdurance Navigator Demo



- Images cycled on old (pre-cycled) blocks to simulate retention period
- Pages containing images are moved from block to block internally
- Every 100 cycles data toggled out
- Images are cycled on default parameter block and also on a Navigator managed block
- 40 bit error detection but no correction. Sectors with uncorrectable errors are deleted

## Summary

- ▶ 3D has made trimming parameters even more difficult
- ▶ Machine Learning is a powerful tool in complex noisy environments
- ▶ FMS 2016 has two commercial deployments of NVMdurance Machine Learning technology
  - ▶ Demonstrating extended life, ultra-flexible deployment
- ▶ NVMdurance Pathfinder is massively scalable
- ▶ Joined up thinking between characterization and deployment is crucial
- ▶ Visit us at Booth 829

