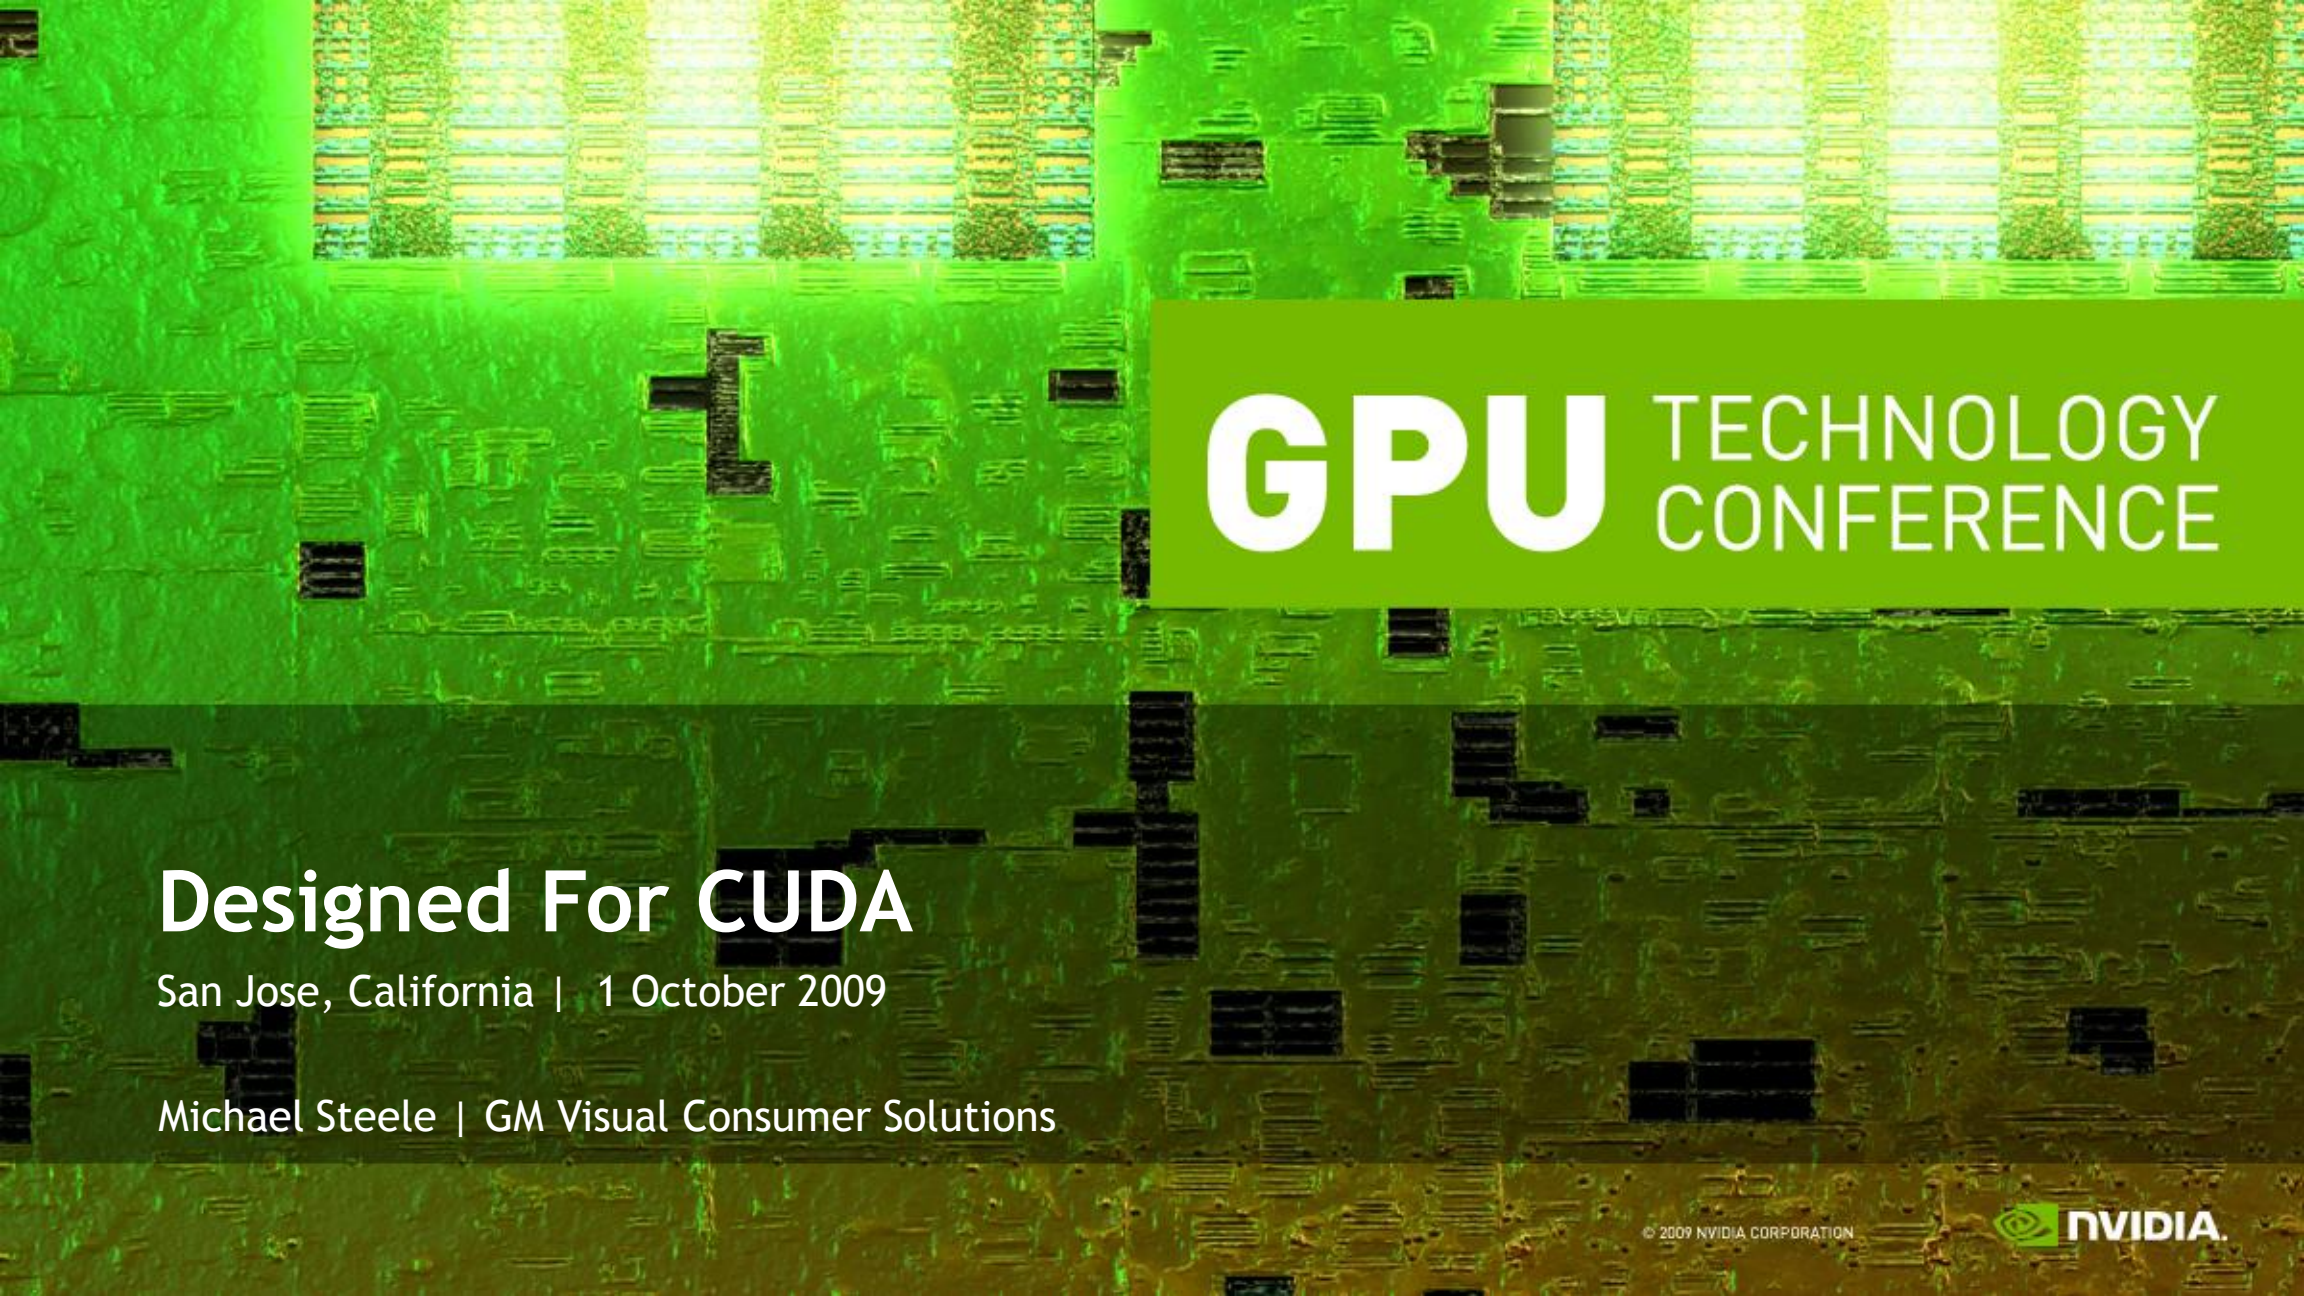

Designed For CUDA

San Jose, California | 1 October 2009

Michael Steele | GM Visual Consumer Solutions

© 2009 NVIDIA CORPORATION

NVIDIA.

Agenda

- The Spirit of “Designed for CUDA”
- How Do You Design for CUDA
- How “Designed for CUDA” Works
  - Development Partnership
  - Marketing Partnership
- Get “Designed for CUDA”

GPU TECHNOLOGY CONFERENCE

DESIGNED FOR NVIDIA CUDA

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

The Spirit of “Designed for CUDA”

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

The Spirit of "Designed for CUDA"

A development and marketing partnership to:

- Enable the best use of GPU computing for your app
- Deliver solutions that amaze consumers
- Build and sustain a profitable market for you and NVIDIA

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

"Designed for CUDA" Program

1 Development and Testing Collaboration For A Better Overall User Experience

- Performance optimizations result in powerful GPU/CPU co-processing
- Stability & reliability via compatibility testing = Fewer returns and support issues
- Advanced GPU computing rewards & encourages new users of CUDA GPUs

2 NVIDIA Market Reach and Application Promotions To Build and Grow Market

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Designed for CUDA - The Approach

CUDA is
An Architecture

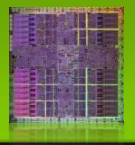

NVIDIA GPU
with the CUDA Parallel Computing Architecture

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Designed for CUDA - The Approach

CUDA is
Open

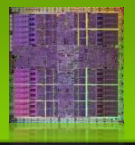

NVIDIA GPU
with the CUDA Parallel Computing Architecture

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Designed for CUDA - The Approach

Widest Language & API Support For GPU Computing

CUDA C	OpenCL	Direct Compute	FORTRAN	Python, Java, .NET, …
• With CUDA Extensions • Over 60,000 developers • SDK + Lib's + Visual Profiler and Debugger	• 1st GPU demo • Shipped 1st OpenCL Conformant Driver • Strategic developers using NV SW today	• Microsoft's GPU Computing API • Supports all CUDA-Architecture GPUs since G80 (DX10 and future DX11 GPUs)	• PGI Compiler • Fortran wrappers • FLAGON	• Compute Kernels • API Bindings

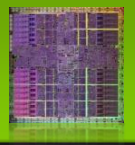

NVIDIA GPU
with the CUDA Parallel Computing Architecture

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Designed for CUDA - The Approach

Your GPU Computing Application

CUDA C	OpenCL	Direct Compute	FORTRAN	Python, Java, .NET, …
• With CUDA Extensions • Over 60,000 developers • SDK + Lib's + Visual Profiler and Debugger	• 1st GPU demo • Shipped 1st OpenCL Conformant Driver • Strategic developers using NV SW today	• Microsoft's GPU Computing API • Supports all CUDA-Architecture GPUs since G80 (DX10 and future DX11 GPUs)	• PGI Compiler • Fortran wrappers • FLAGON	• Compute Kernels • API Bindings

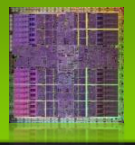

NVIDIA GPU
with the CUDA Parallel Computing Architecture

© 2009 NVIDIA Corporation

NVIDIA

Designed for CUDA - The Approach

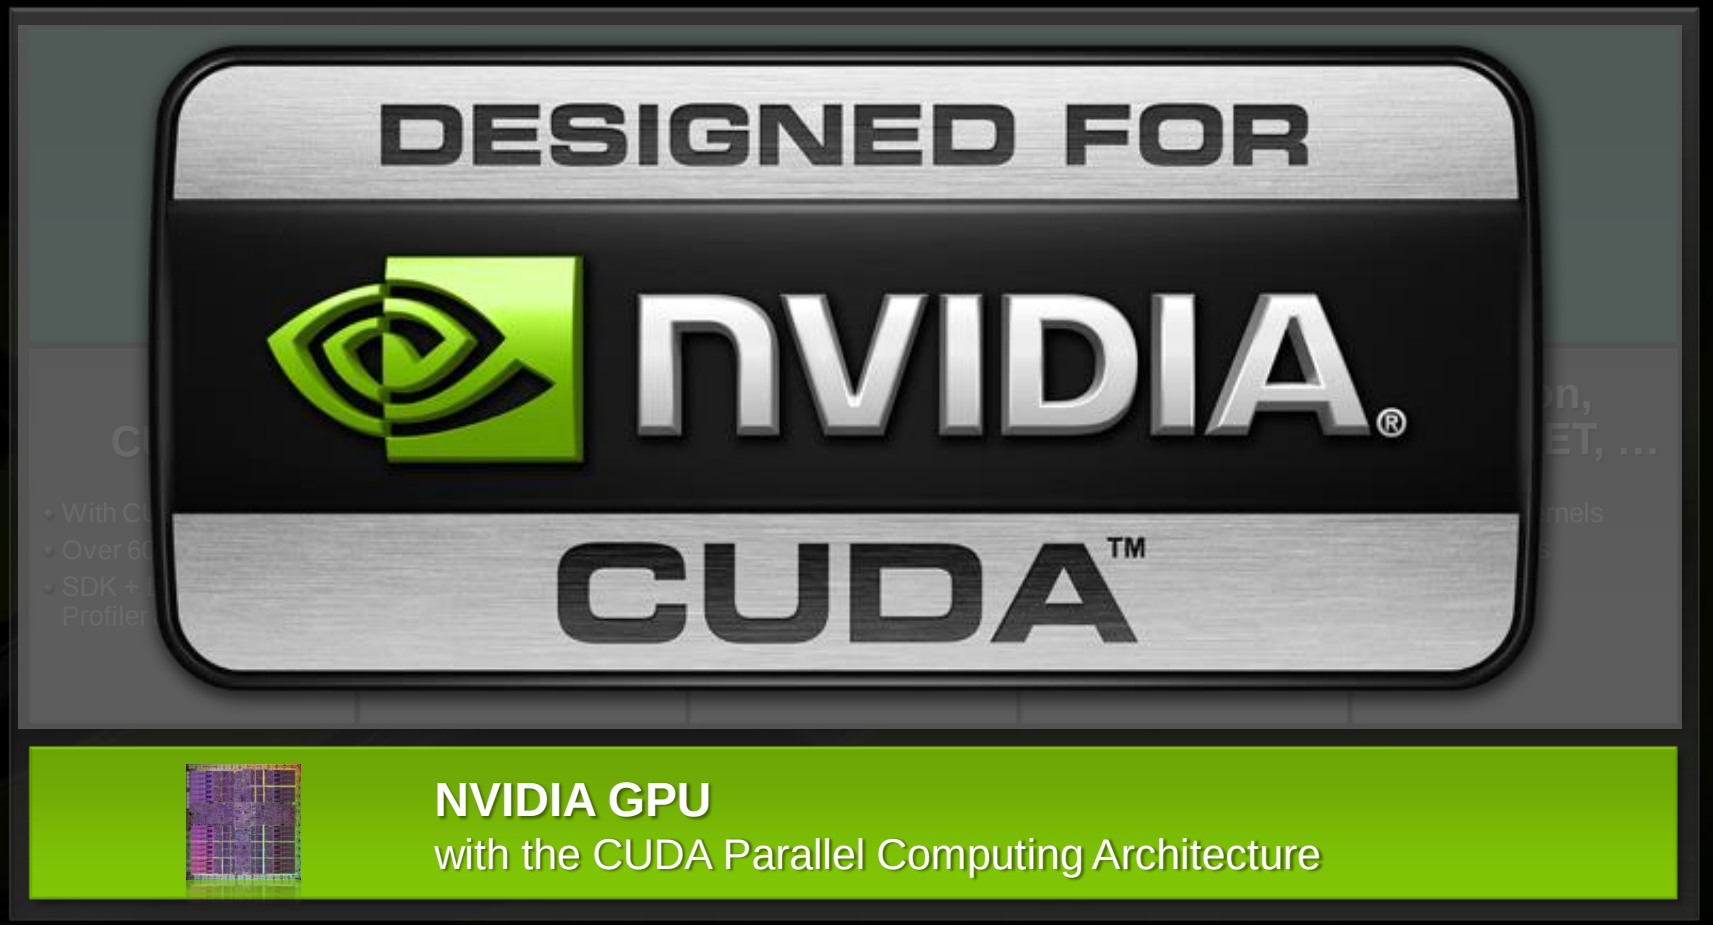

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Designed for CUDA - The GPU

Over 100 Million CUDA-enabled GPUs Installed

GeForce®
Consumer

Tesla™
HPC | Server

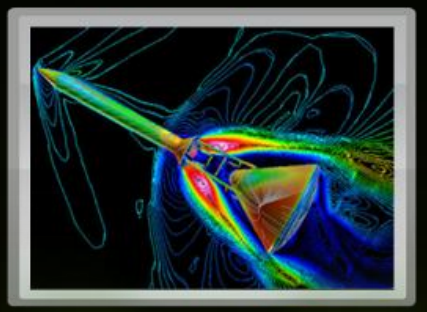

Quadro®
Professional

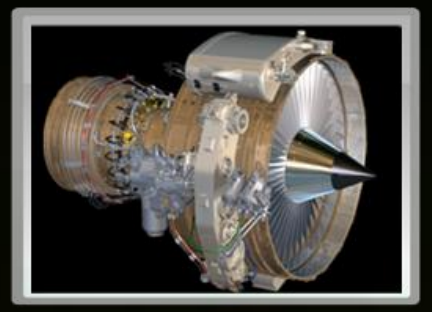

© 2009 NVIDIA Corporation

NVIDIA.

GPU TECHNOLOGY CONFERENCE

Designed for CUDA
Reaching Consumers

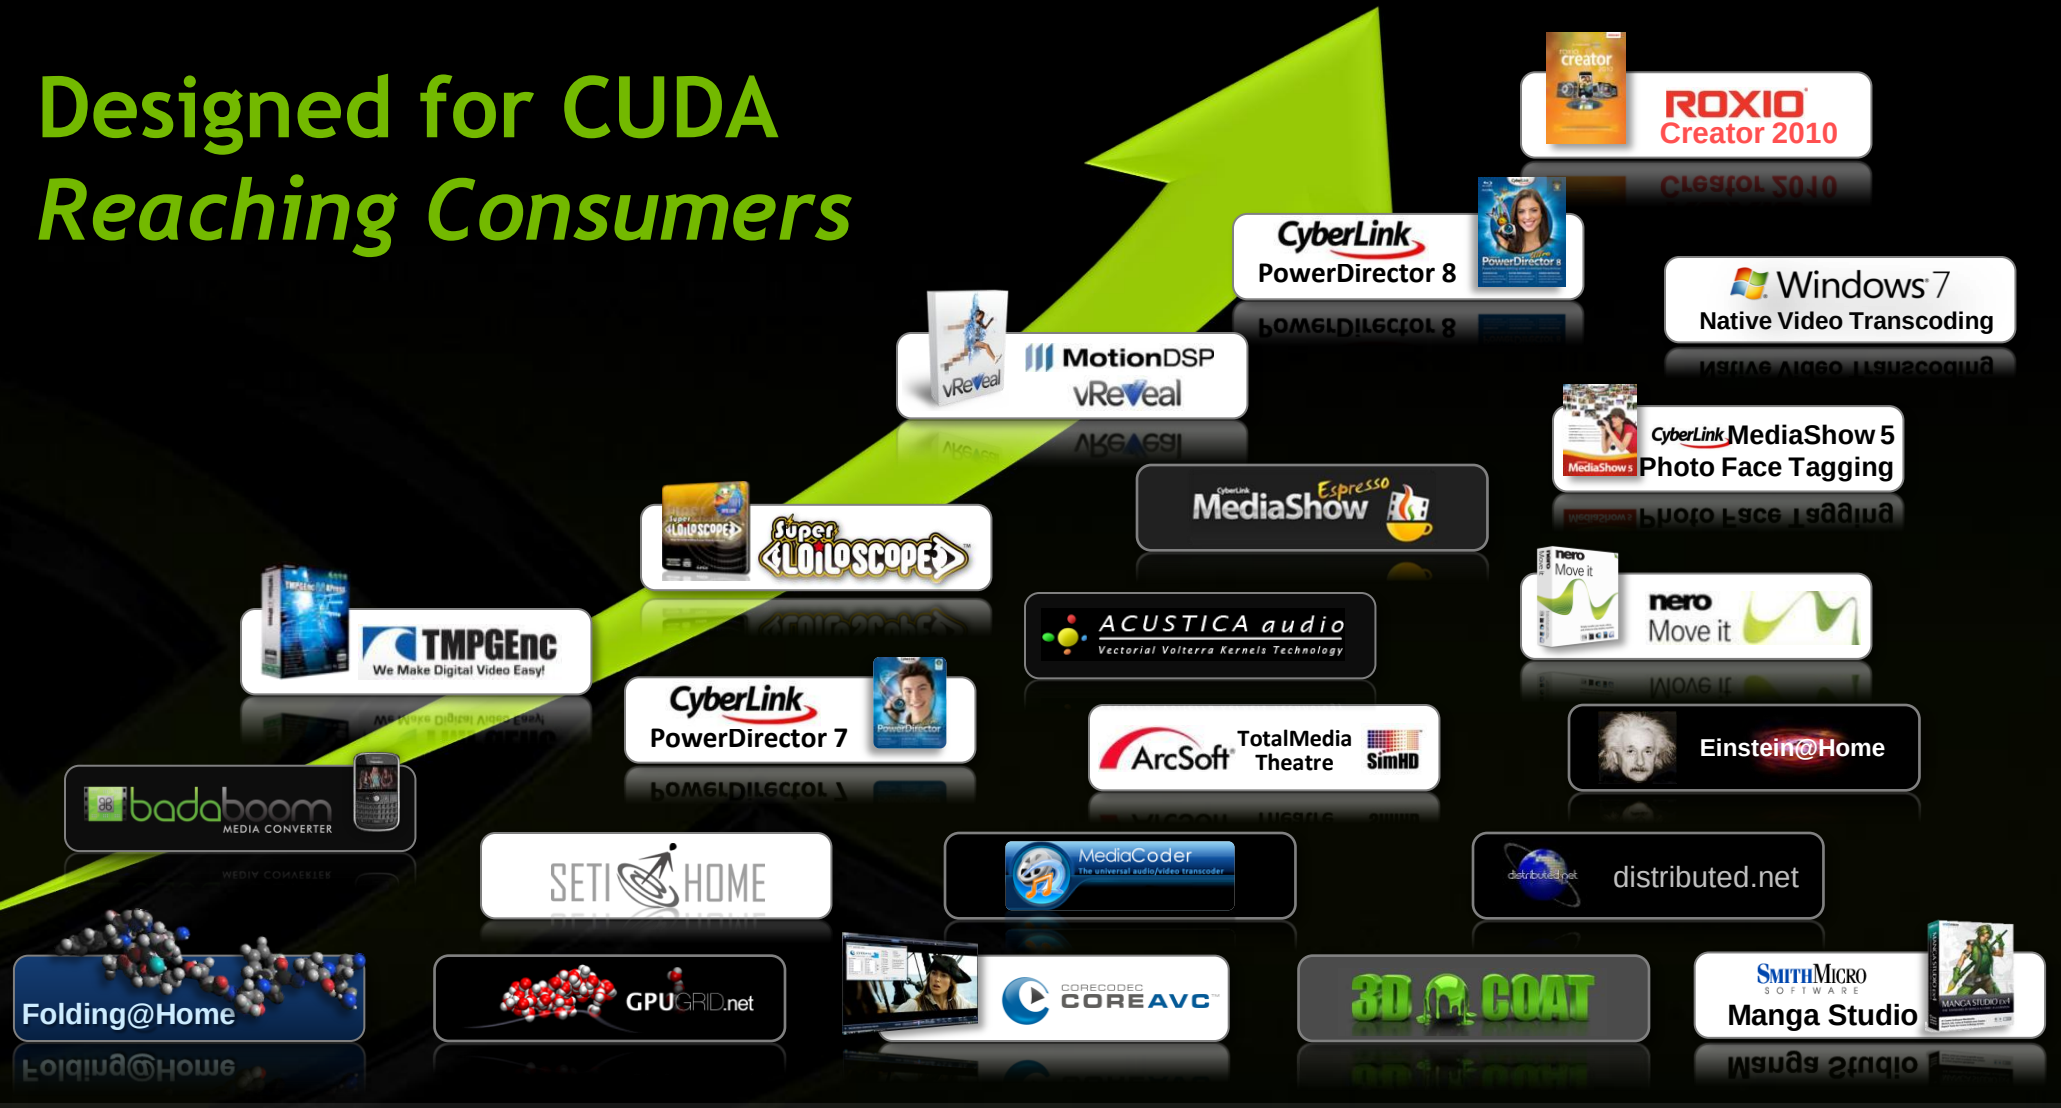

2H 08 | 1H 09 | 2H 09

© 2009 NVIDIA Corporation

NVIDIA

Program In A Nutshell

NVIDIA Provides to ISV Partner

- **Technical info**
  - Early access to future products, tech details & roadmaps
- **Development support**
  - Implementation, optimization, bug fixes to best enable NVIDIA GPU
- **Testing support**
  - Compatibility testing, min spec analysis
- **Marketing programs**
  - Targeting 100M+ GPUs

ISV Partner Provides to NVIDIA

- **High impact CUDA technology adoption**
- **App info**
  - Tech details, release & update timing
- **Early and regular access to builds & updates**
- **Ongoing GPU computing optimization**
- **Co-marketing collaboration**

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

NVIDIA Value Proposition — Development

- **Extensive Technology and Tools**
  - Optimized GPU Computing support
  - Expertise in Video, Imaging, & GPU Computing
  - Advanced graphics feature support
  - Performance optimizations
- **Global Test Lab with Over 800 Systems**
  - Dedicated to pre-release app testing
    Santa Clara, Moscow, Shanghai, Beijing
  - GPU scaling analysis
- **Inside Track Access**
  - NVIDIA developer website [*private*]
  - File and track bugs
  - Early access: drivers, tools, roadmaps

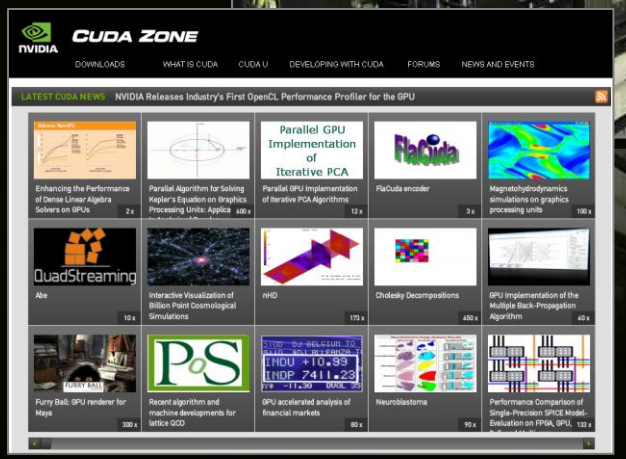

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

NVIDIA Developer Resources

DEVELOPMENT TOOLS

CUDA Toolkit
Complete GPU computing development kit

Visual Profiler
GPU hardware profiler for CUDA C and OpenCL

Nexus
Development environment with Visual Studio integration [*beta*]

NVPerfKit
OpenGL|D3D performance tools

FX Composer
Shader Authoring IDE

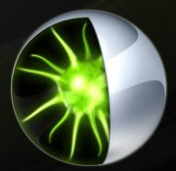

SDKs AND CODE SAMPLES

GPU Computing SDK
CUDA C, OpenCL, DirectCompute programming guide and code samples

Graphics SDK
DirectX & OpenGL code samples

PhysX SDK
Complete game physics solution

OpenAutomate
Test automation SDK

VIDEO LIBRARIES

Video Decode Acceleration
NVCUVID
DXVA
Win7 MFT

Video Encode Acceleration
NVCUVENC
Win7 MFT

Post Processing
Noise reduction / De-interlace/ Polyphase scaling / Color process

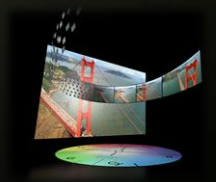

ENGINES & LIBRARIES

NPP Image Libraries
Performance primitives for imaging

Numeric Libraries
cuFFT, cuLA, cuBLAS

App Acceleration Engines
Optimized software modules for GPU acceleration

Shader Library
Shader and post processing

Optimization Guides
GPU computing and Graphics development best practices

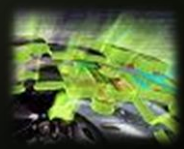

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Helping You Reach Over 100 Million Consumers

Source: Mercury Research, June 2009

© 2009 NVIDIA Corporation

NVIDIA

GPU TECHNOLOGY CONFERENCE

Leverage our Global Marketing Reach

nvidia.com

Driver Page

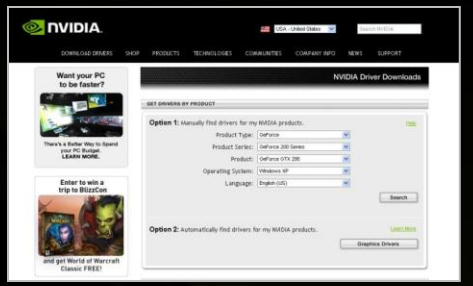

nZone.com

OEM & Card Bundles & Mkg

Apple, ASUS, SAMSUNG, SONY, LG, hp, IBM, DELL, XFX, BFG Tech, PNY TECHNOLOGIES, EVGA

PartnerForce

WW Channel Partners

High Impact PR

"GPU computing may be coming into its own, with the first real consumer-ready apps." — PC Magazine

"… power of parallel computing is amazing." — TG Daily

"… game changing technology" — CPU Computer Power User

"… advantages of GPU-equipped systems will eventually become overwhelming. . . no longer recommend PCs without GPUs." — cnet news

Trade Shows | Events

NVISION Magazine

© 2009 NVIDIA Corporation

NVIDIA

Co-Marketing Assets

- Banner ad templates
- Print ad templates
- Mouse pad
- Monitor hats
- Shelf talkers
- Poster
- Box flap stickers

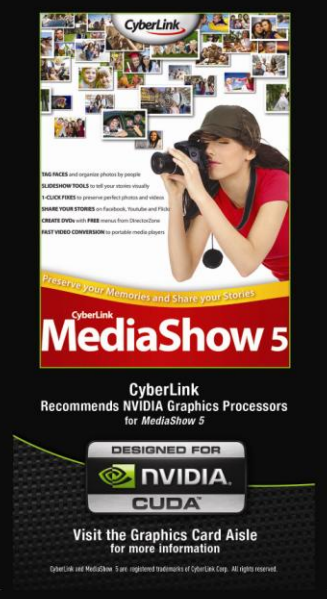

© 2009 NVIDIA Corporation

Leverage NVIDIA Partner Connections

- Leading OEM's [Desktop & Notebook]
- Over 25 add-in card partners (AIC's)
- 1,000s of resellers via "Partner Force" program
- Distributors
- Regional System Builders

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

Co-Market with NVIDIA Partners!

- Promo ads / flyers
- Web banner ads
- Shelf talkers
- Box promo

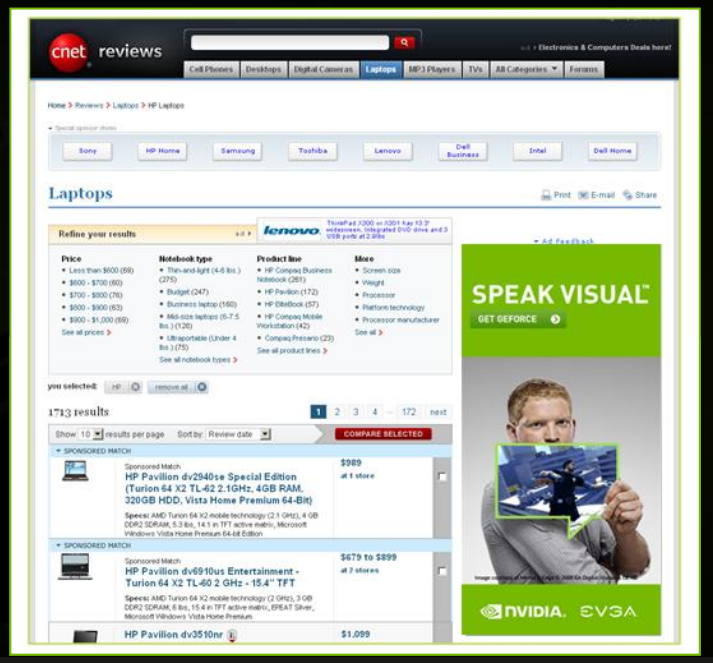

© 2009 NVIDIA Corporation

It's Your Move

- Contact NVIDIA Now
  - Email - designedforcuda@nvidia.com

- Develop
  - Enable the best experience
  - Optimize to expand your market
  - Ensure stability with advanced testing

- Market
  - Reach over 100+ million CUDA users
  - Tap NVIDIA partners to increase visibility
  - Sell via NVIDIA online stores

© 2009 NVIDIA Corporation

GPU TECHNOLOGY CONFERENCE

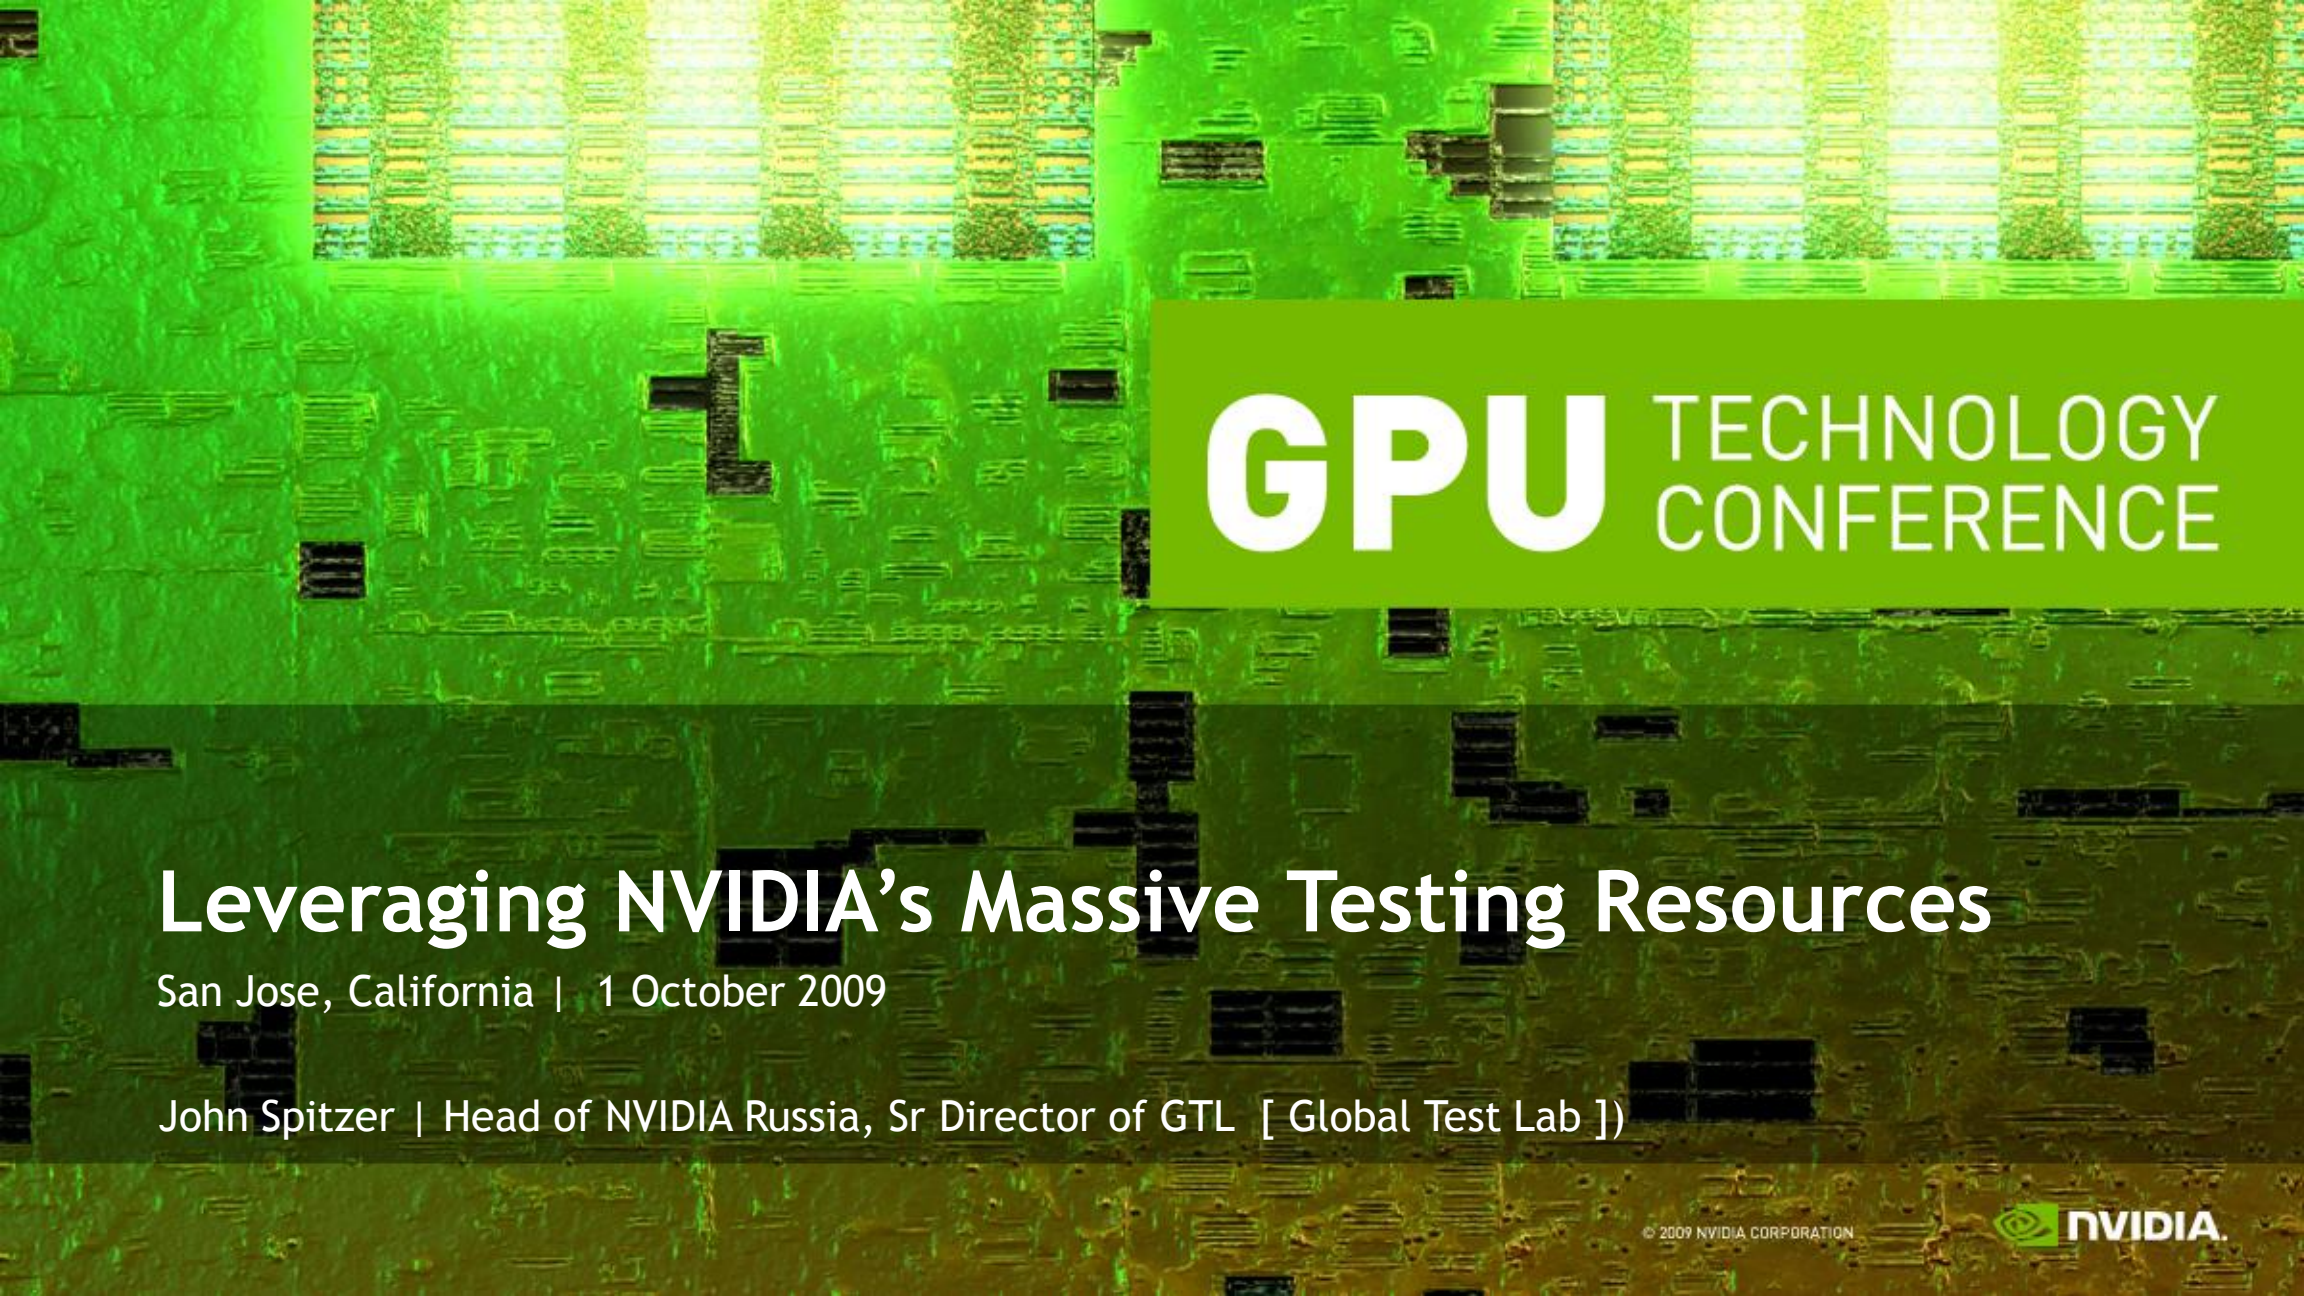

Leveraging NVIDIA's Massive Testing Resources

San Jose, California | 1 October 2009

John Spitzer | Head of NVIDIA Russia, Sr Director of GTL [Global Test Lab])

© 2009 NVIDIA CORPORATION

NVIDIA.

Why Do We Need Automation?

- People don’t scale as well as machines
- Reproducibility
- Reduces potential for error
- Faster turnaround time

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

OpenAutomate

- Purpose
  - Automation for QA and benchmarking
  - Standard for querying/setting application parameters
  - Standard for running automated benchmarks
- Design goals
  - Minimal API, DLL, data structures
  - No complex file formats to parse/write
  - Virtually no footprint added to application

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

Benefits

Reduces customer support issues

- Manual testing allows compatibility testing on just a handful of configurations
- Automation allows testing on the full spectrum of configs in NVIDIA's lab (>500)
- Not just GPU, but CPU, memory, display, OS, etc..

Forward compatibility

- Top-to-bottom testing on pre-release drivers, compilers, runtimes and so on
- Insurance that your game will run well on future GPUs

Comprehensive reports

- Compatibility, but also…
- Suggested minimum specification
- Performance data

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

Benchmark Requirements

- Type and length
  - Key word is “representative” - meaning real-life
  - Too short (less than 30s) and timing error can become significant
  - Too long (more than 15 min) will bog down our farm
  - Many short benchmarks (3 min) is often better than one long (15 min)
- Deterministic benchmarks are ideal
  - 2 runs, 1 machine, identical results
  - Even better: 2 runs, 2 machines, identical results

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

How Does It Work?

- The app is invoked with a command-line arg
- The app call oaInit() with the command-line option
- oaInit() loads a DLL, and binds all the oaXXX funcs
- The app goes into a command-loop, waiting for commands from the DLL

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

What Are The Possible Commands?

- Enumerate available options
- Get current options
- Set options
- Enumerate available benchmarks
- Run benchmark
- Run application
- Exit

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

Basic Implementation:

Application

Command-line: app.exe –openautomate OA_plugin.dll;

1. Takes OpenAutomate command-line and starts initialization of OA:

```
oaVersion Version;
oaInit((const oaString)OAOpt, &Version);
```

3. Enters OA command loop:

```
oaInitCommand(&Command);
switch(oaGetNextCommand(&Command))
```

6. Once OA_CMD_EXIT command is received, the application exits

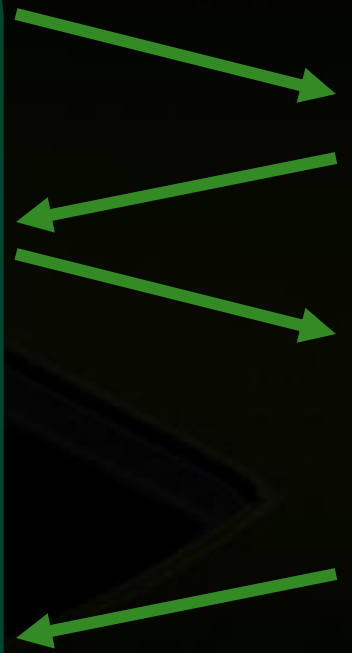

OA_plugin.dll

2. Initializes and returns function pointers for callbacks

4. DLL sends OA command stream to app:

OA_CMD_GET_ALL_OPTION
OA_CMD_GET_CURRENT_OPTIONS
OA_CMD_SET_OPTIONS
OA_CMD_GET_BENCHMARK
OA_CMD_RUN_BENCHMARK
OA_CMD_RUN

5. Eventually the DLL will send the OA_CMD_EXIT command.

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA.

Submission Checklist

- ❑ No installation (‘copy and run’), silent if former isn’t possible
- ❑ Provide the entry point exe name
- ❑ No licensing needed (or if necessary, unlimited licenses with clear instructions on how to do so)
- ❑ Full OA conformance (with conformance log)
- ❑ At least 1 benchmark exposed in OA, no longer than 15 minutes in length
- ❑ No user interaction should be required while in OA mode
- ❑ Barring major errors, the application should never exit by itself unless OA_CMD_EXIT is received
- ❑ Make sure to call oaStartBenchmark and oaEndBenchmark in appropriate places

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA

Call to Action

- Increase your testing coverage by 100x
- Insure compatibility with future drivers and GPUs
- It's easy to implement
- OpenAutomate is now public! http://openautomate.com
- E-mail: OpenAutomate-Support@nvidia.com

GPU TECHNOLOGY CONFERENCE

© 2009 NVIDIA Corporation

NVIDIA.

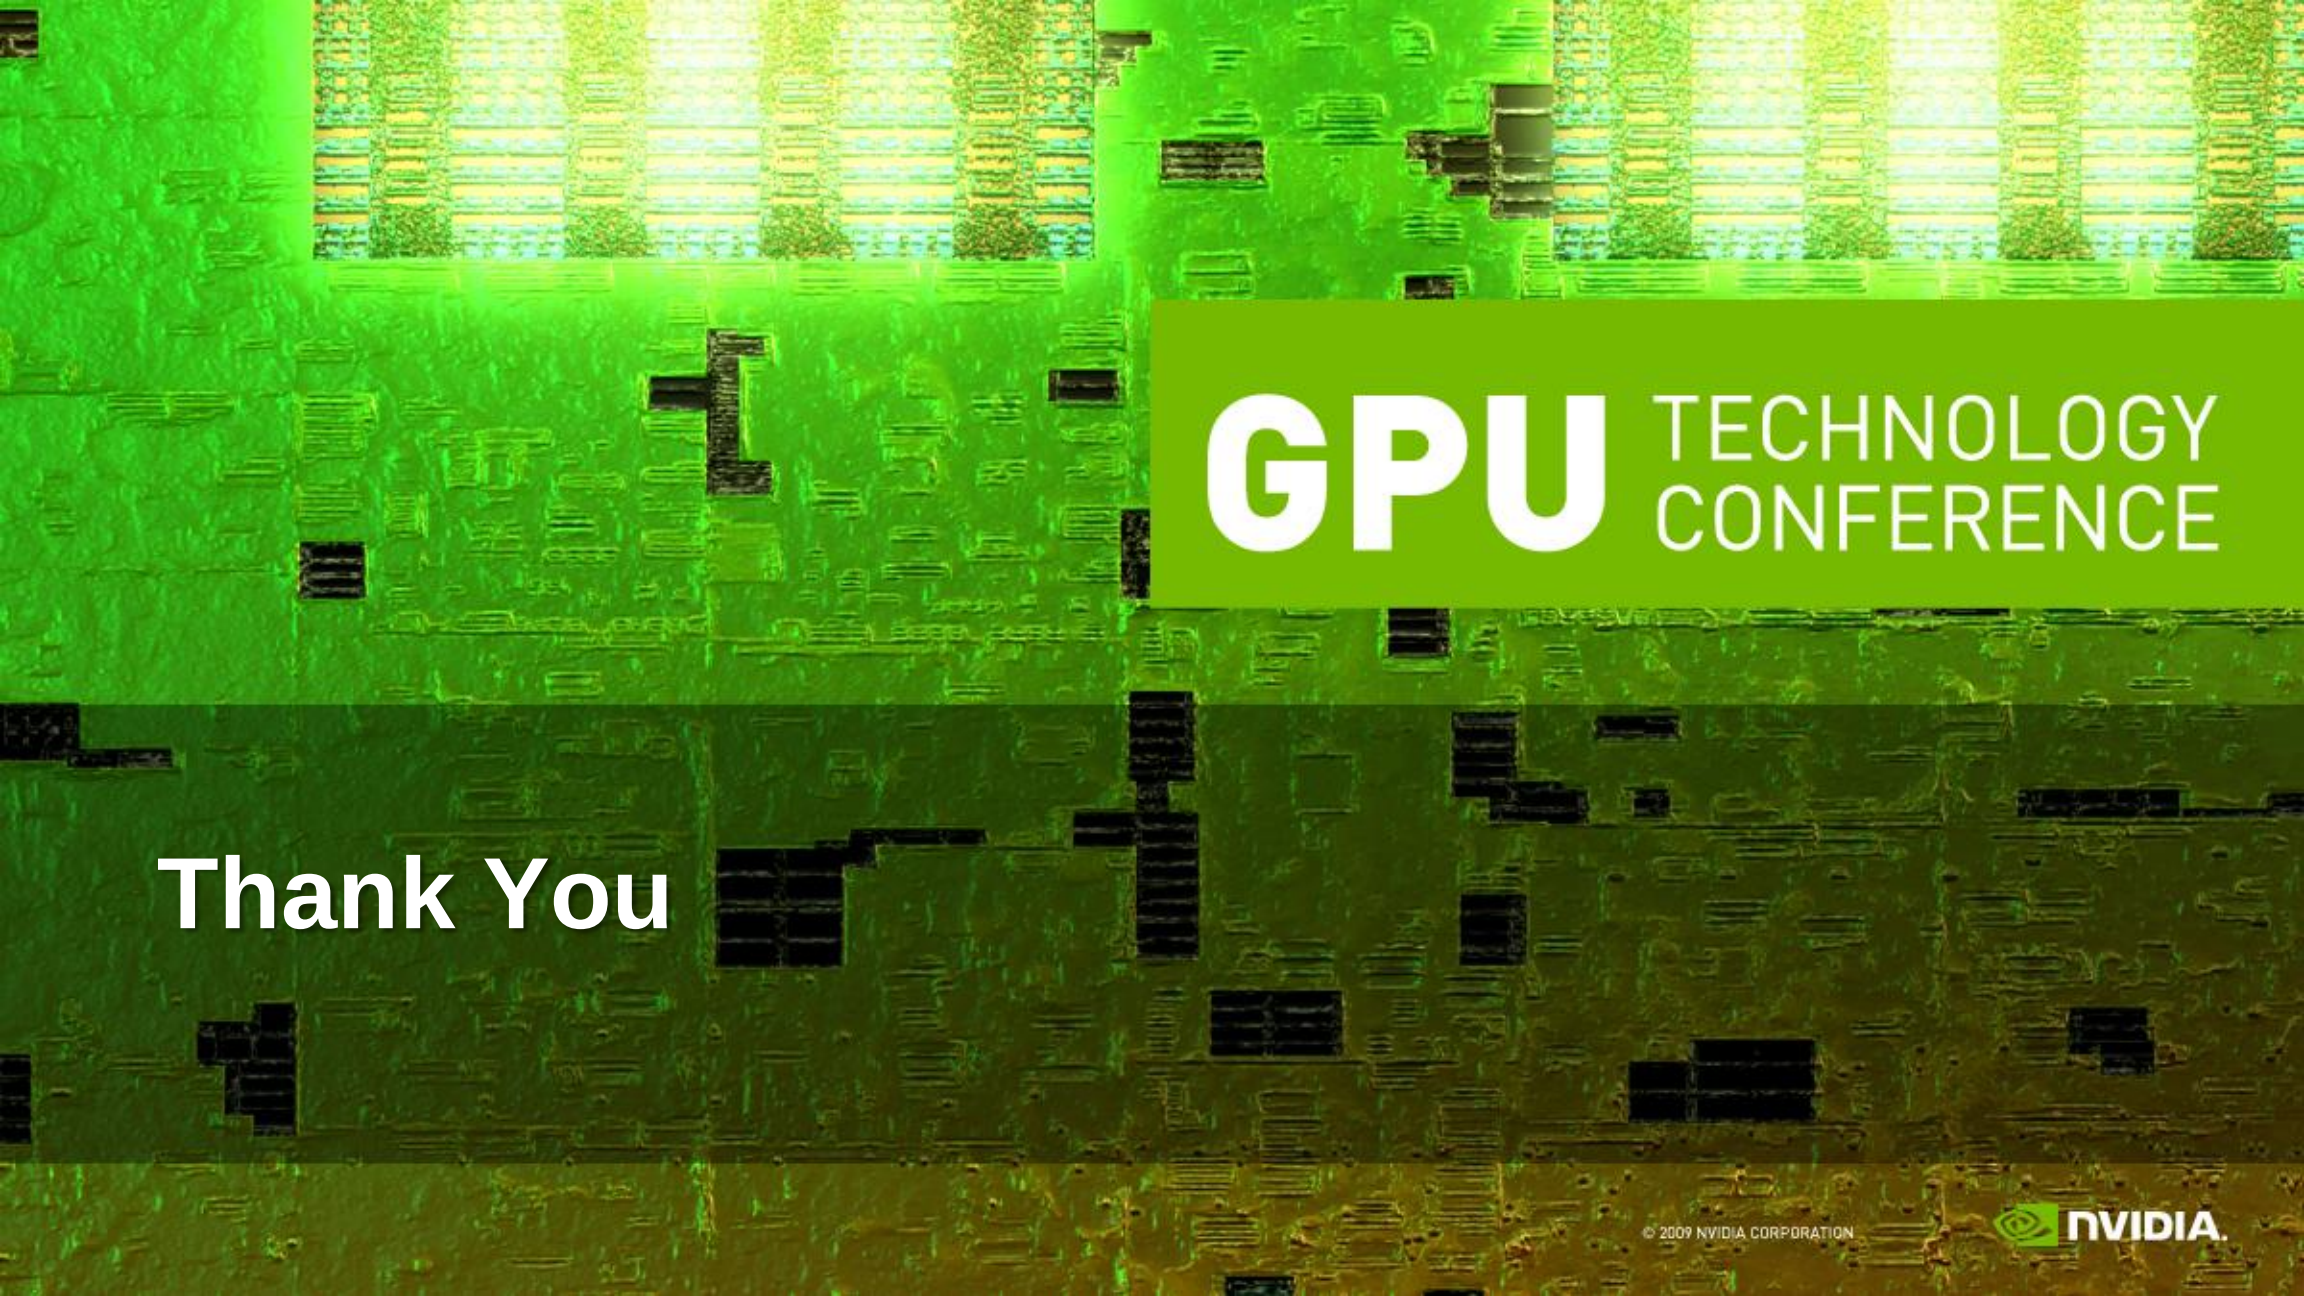
GPU TECHNOLOGY CONFERENCE
Thank You
© 2009 NVIDIA CORPORATION
NVIDIA.