TEXTRON

Exploring Design Considerations

CAD/CAM Experiences from the Experts using Citrix and VMware

NVIDIA GTC 2015

Fred Devoir & Randall Siggers

Who are we?

Fred Devoir

Sr. Architect
TEXTRON

MIS, EMCCAe, ITILv3

Relevant VDI Experience:

- **Citrix VDI solution for NASA International Space Station**
- **Citrix HDX 3DPro development team member on customer side NASA/Boeing**
- **Citrix XenDesktop 6.5 GPU pass-through solution for NASA Engineering Directorate at JSC.**
- **Citrix XenDesktop 7.6 GRID K2 enabled vGPU solution for Bell Helicopter.**

Randal Siggers

Solutions Architect
TEXTRON

Relevant VDI Experience:

- **GRID enabled vSGA VMware solution for Jacobs Engineering**
- **Working with Nvidia GRID team to validate vGPU for VMware**
- **Working on GRID enabled VMware vGPU POC for TEXTRON Tool and Test**

TEXTRON

Disruptive Technology Paradigm

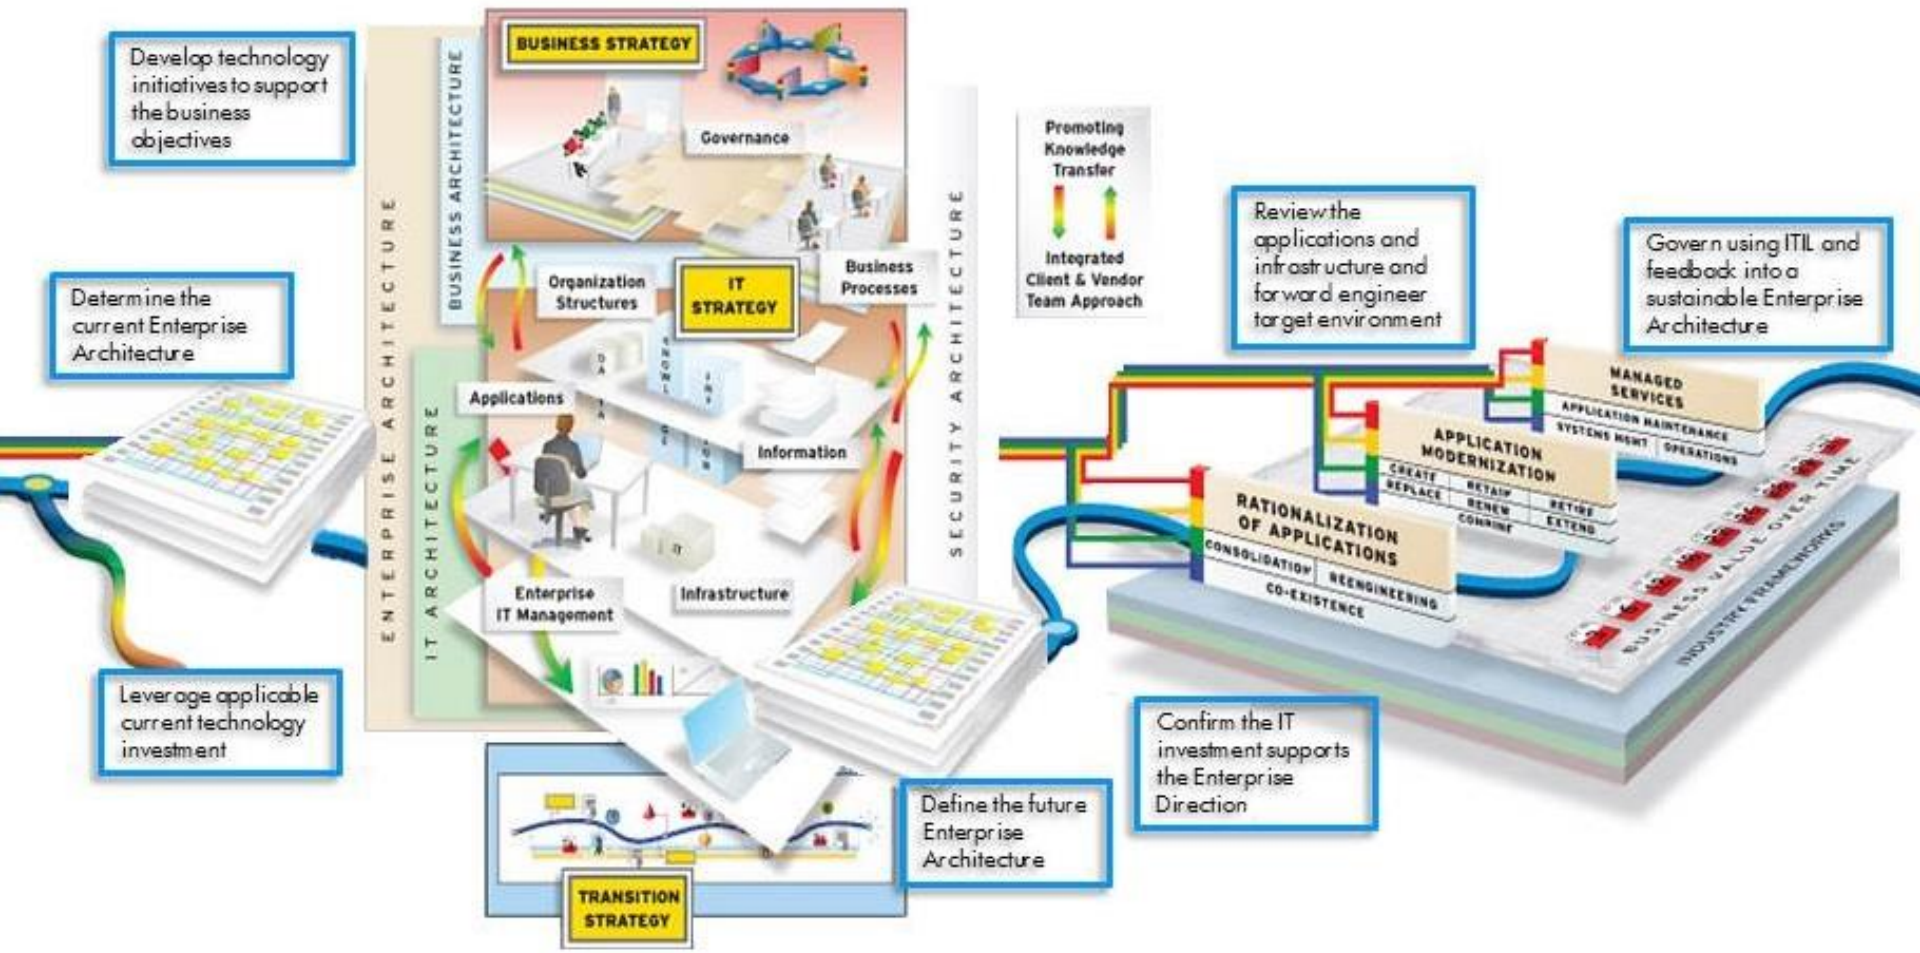

TEXTRON

Traditional CAD/CAM Challenges

- **Expensive machines and workflow restricted to a local environment**
- **Multiple physical assets to accommodate varying project ecosystems**
- **Limited ability to work share beyond a single location**
- **Disparate design teams used software collaboration tools to circumvent the fact that the design files, compute and GPU resources were not collocated**

Transitional Design Challenges

- Politics – psychology vs technology
- Distributed compute and data
- Complex support model {N:*f(users)*}
- Application diversity
- Device creep – endpoints and operating system images

Solutions Provided by VDI

- Centralized compute and data freedom
- Simplified support model {$f(techs)$:N}
- Application standardization and version control
- Operating system control with endpoint ambiguity

Design Observations

- Rack density depends on hardware choice
- Power & cooling depends on hardware choice
- User density depends on vGPU profile
- Cost neutrality

The Apex Predator vs Mammals

	PERSONAS	ROLE DEFINITION	ATTRIBUTES	EXAMPLES
Mammals	**Small VDI (White)** **2 vCPU 16GB RAM 1GB vGPU 96GB wCache 15,000 IOPS**	Standard User	▪ View-only application user accessing individual parts or small assemblies ▪ Full application user accessing individual parts or small assemblies ▪ Requires less than 1 GB of GPU Memory ▪ Requires 16 GB or less of RAM	PLANNERS SCHEDULERS QUALITY CONTROL SHOP FLOOR ENG
Apex Predators	**Medium VDI (Red)** **4vCPU 32GB RAM 1GB vGPU 152GB wCache 15,000 IOPS**	Moderate User	▪ View-only application user accessing medium assemblies ▪ Full application user accessing medium assemblies ▪ Requires 1 GB of GPU Memory ▪ Requires 16-32 GB of RAM	ENG MANAGEMENT AVERAGE ENG USER BUSINESS PARTNER
	Large VDI (Blue) **4 vCPU 64GB RAM 2GB vGPU 224GB wCache 15,000 IOPS**	Power User	▪ Application user accessing large assemblies or full aircraft ▪ Requires 2 GB of GPU Memory ▪ Requires 32-64 GB of RAM	DESIGN ENG ENG MANAGEMENT BUSINESS PARTNER
	Jumbo VDI (Purple) **4 vCPU 128GB RAM 4GB vGPU 768GB wCache 15,000 IOPS**	High End Power User	▪ Application user accessing large assemblies or full aircraft ▪ Requires 4 GB of GPU Memory ▪ Requires 64-128 GB of RAM	DESIGN ENG ANALYSIS ENG

TEXTRON

Ali Rizvi

PLM Support Analyst
Bell Helicopter

Email: srizvi01@bh.com

Twitter: @alirizvi022

LinkedIn:

http://www.linkedin.com/in/syedalirizvi1

TECHNOLOGY DEMO

TEXTRON

Bell Helicopter – Production Environment

- **VDI 112 User Production Deployment**
  - The VDI production Phase 1 equipment is 16 Lenovo nx360M4 blades

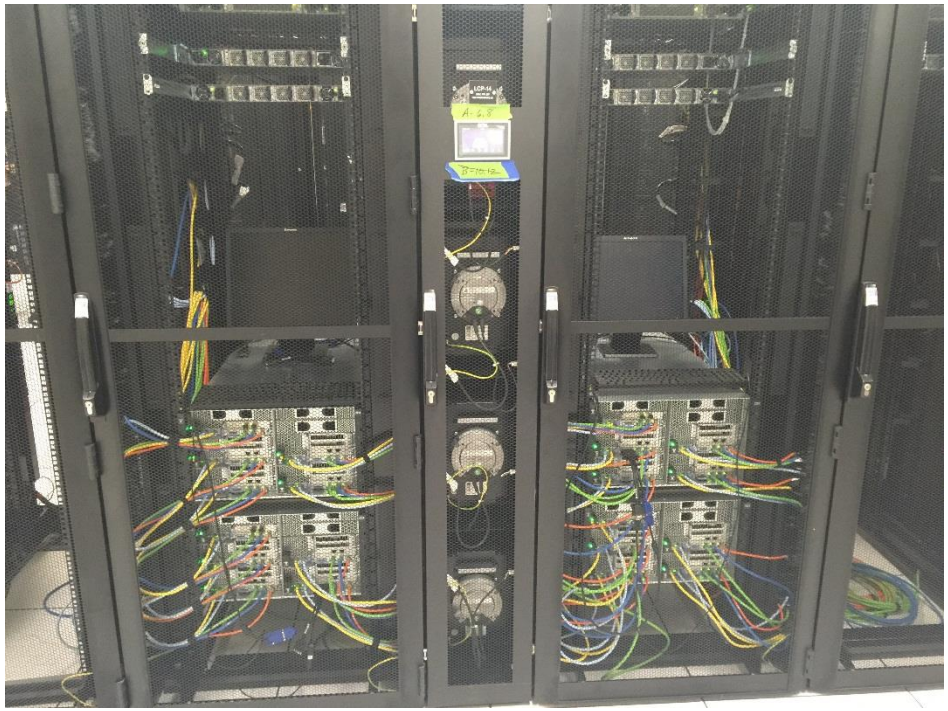

- **Each blade is equipped with dual Nvidia GRID K2 GPUs**
- **Each blade has 2x 2.2Ghz 10Core processors, 256GB RAM, and 4x 800GB SSD disks (3.2TB) for write cache.**
- **The blade are powered by 6x 1300Watt power supplies per chassis with 4 blades per chassis.**
- **Each blade is connected with dual 10GB network uplinks.**

PoC Environment

- **VDI PoC for VMware, Citrix, and HyperV**

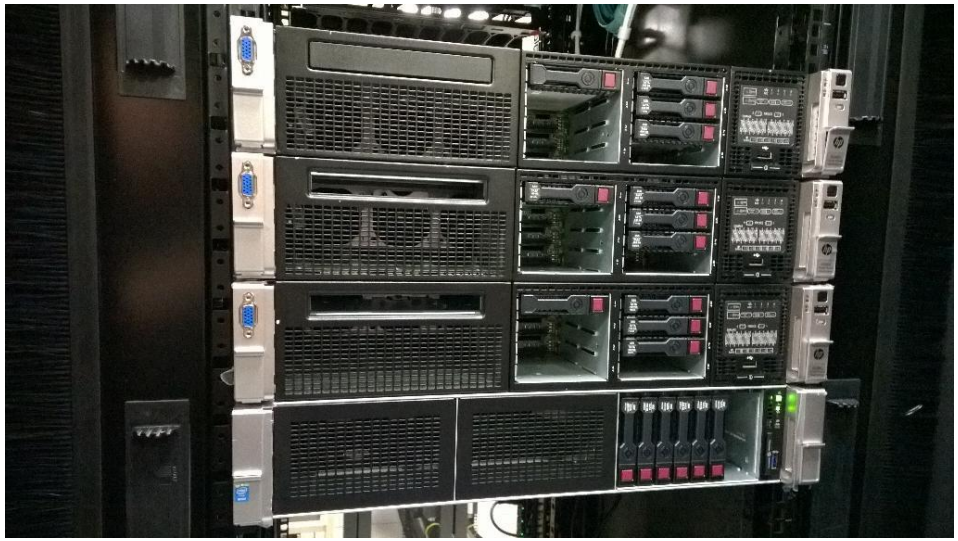

- **3x HP DL380z Gen8, 1 HP DL380z Gen9 each equipped with dual Nvidia GRID K2 GPUs**
- **Each server has 2x 3.00GHz 10Core processors, 384-512GB RAM, and 3x 480GB SSD disks (900GB) for write cache.**
- **The server are powered by 2x 1400Watt power supplies**
- **Each server is connected with single 1GB network uplink (because its PoC)**

QUESTIONS

TEXTRON

Contact Us

Fred Devoir

Sr. Architect
TEXTRON

MIS, EMCCAe, ITILv3

Email: fdevoir@textron.com

Twitter: @devoirf

LinkedIn:

http://www.linkedin.com/in/devoir

Randal Siggers

Solutions Architect
TEXTRON

Email: rsiggers@textron.com

Twitter: @siggtonium

LinkedIn:

http://www.linkedin.com/in/rsiggers

Thanks and Acknowledgements

CITRIX® IBM NVIDIA® hp vmware®

lenovo Microsoft CISCO AppSense ARISTA