

STORAGE 101:

PAST, PRESENT, FUTURE OF STORAGE.



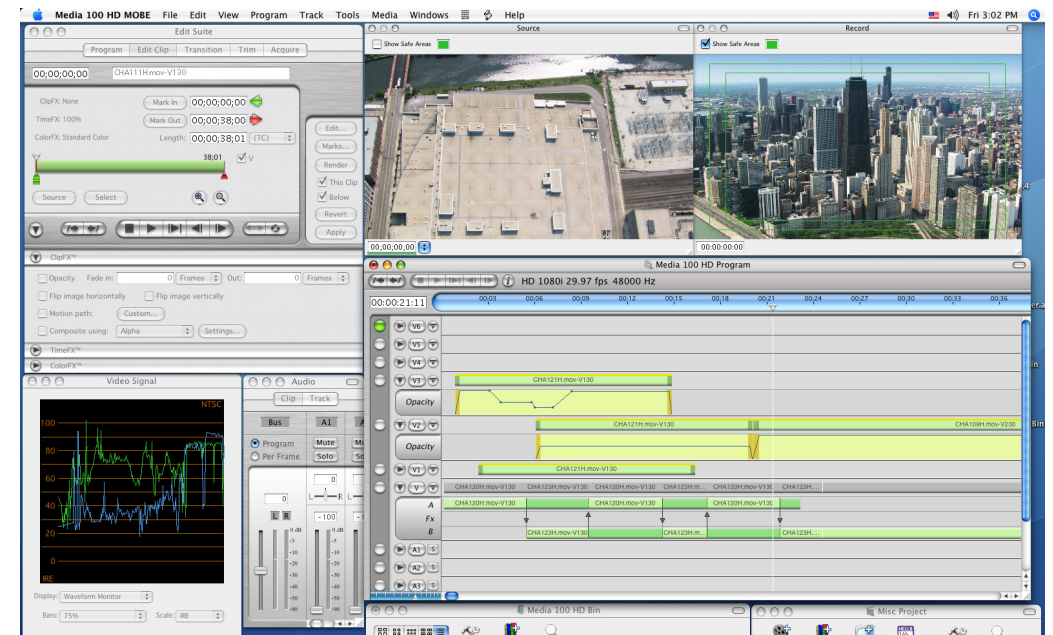
creative.space[®]
powered by DigitalGlue #

HARD DRIVE TECHNOLOGIES

- **PATA DRIVES** Parallel ATA (Parallel Advanced Technology Attachment or **PATA**)
- **SATA DRIVES** Serial ATA (Serial Advanced Technology Attachment or **SATA**)
- **SAS DRIVES** Serial Attached SCSI (**SAS**)
- **FATA DRIVES** Fibre Attached Technology Adapted (**FATA**)
- **SSD** Solid-State Drive (**SSD**)
- **NVMe** Non-Volatile Memory Express (**NVMe**)

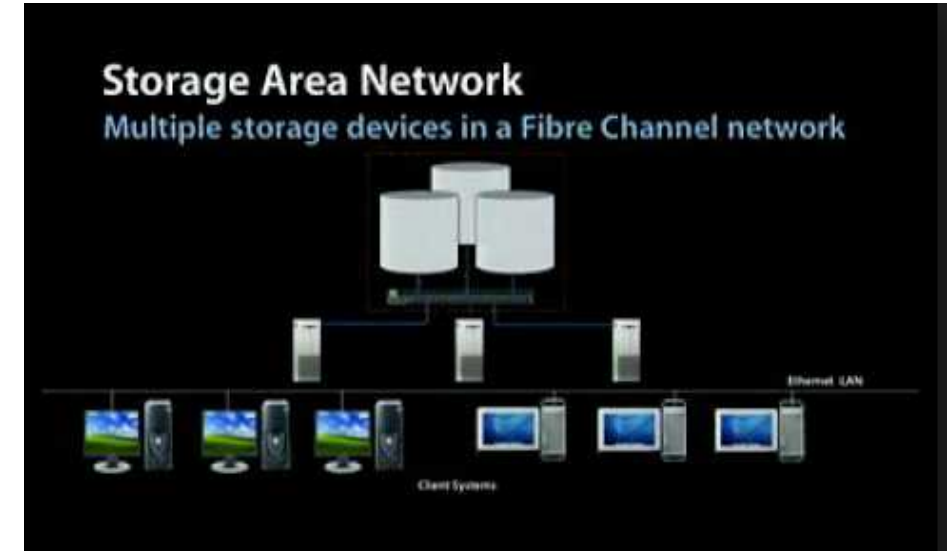
GETTING VIDEO TO THE DESKTOP

- In the 90's Lockheed Martin was looking to bring video to the desktop.
- What technology did they use?
 - Fibre Channel
- What went wrong?
 - File churn that created lots of extents
- What evolved from the project?
 - Stornext



KEEPING VIDEO TO THE DESKTOP NAS OR SAN

- Starting around the year 2007 post productions
 - studios needed to make a decision NAS or SAN
- Why did we need to use a SAN?
 - Many decided with the SAN as it offered speeds around 400 MB/sec
- Why not use a NAS?
 - Ethernet only offered 1GbE
- What are the disadvantages of a SAN?
 - Difficult to DPX, Fragmentation, many failure points



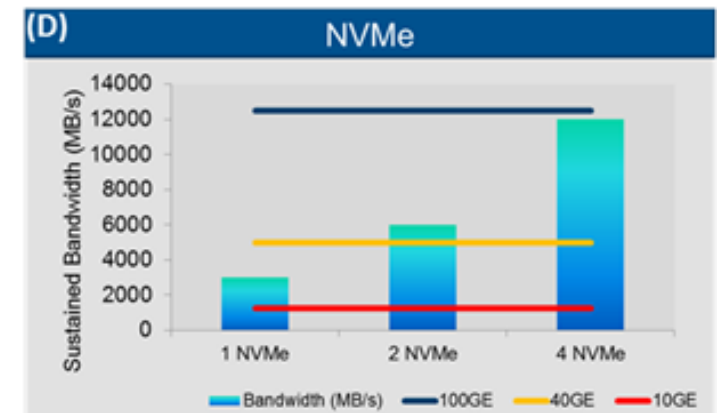

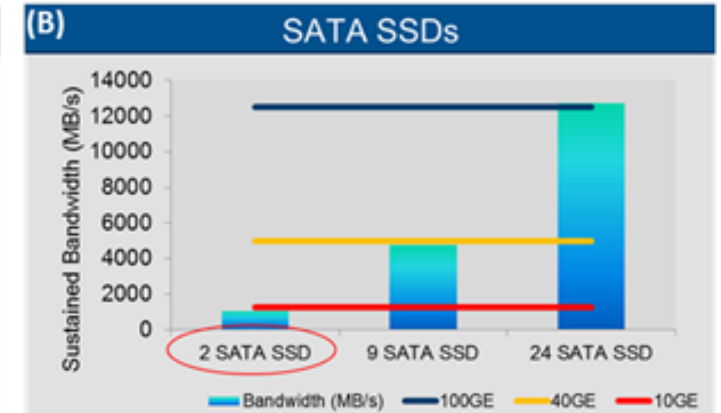

NVMe vs. SSD



- SSD debuted in 2009
- What's held SSD back?
 - 6 Gb/SATA (550MB) cable or 12Gb/SAS (1,100MB)
- Why is NVMe so much faster?
 - Connects directly to the PCIe3 X4 bus
- How will we use NVMe in M&E?
 - 2TB delivering 3,000 MB/sec
- Will NVMe be the future?
 - No, it's the present

THE NUMBERS


- How many drives to fill a 100GbE pipe?
 - SATA HDDs – 250
 - SATA SSDs – 24
 - SAS SSDs – 10
 - NVMe – 4



WHAT IF SOMEONE WENT EXTREME?



- 36 NVMe Drives in 1 RU
- 10 Millions IOPS
- 576 TB



HOW DO WE LEVEREDGE NVMe?



- The technology was first defined by the IEEE 802.3ba-2010 standard
- 40GbE and 100GbE made their debut
- What can we do with this technology
 - 100GbE from NAS to 10GbE Switch
 - Support Image Sequences (DPX, TIFF, OpenEXR) and demanding bitrate(s)

THANK YOU

STEVE GRAPPONE
SYSTEMS ARCHITECT, DIGITAL GLUE
steve.grappone@digitalglue.com

creative.space[®]

powered by DigitalGlue 

creative.space

powered by DigitalGlue 



Mellanox – Ethernet Storage Fabric

Presented by Bill Webb, Director - Ethernet Switching - Americas

September 17, 2018






Mellanox Ethernet Switches Power All Platforms

 **3,000+** unique customers worldwide

 **Over 1M** ports deployed

 **OEM'd** by global server OEMs

 **Open** Network Operating Systems



Mellanox Accelerates Businesses' Success!

Reducing CPU Usage Drastically with Advanced Network Offloads
IP-Broadcasting Applications and Many More on 25/100GbE!



100K subscribers from a Single Caching Server!
Netflix Open Connect Appliance with Mellanox ConnectX-4



Mellanox 40GbE-based Network to Support Clustered File Systems
Live Broadcast Video Editing



Mellanox 10/40GbE-based Storage Network
Offline Broadcast Editing




Only End-to-End 100GbE Network to Support Maximum of 32 Tiles
Enabling IP-Broadcast Multiviewer & Signal Analytics



25G & Above is Fastest Growing Ethernet Segment


Mellanox is #1 with 67% share in the Nx25G segment

High Speed Adapters
2017 (\$1.3B)




22%
Of Total
Market

25G+ Segment Growing at 38% CAGR



10G 25G 40G 50G 100G 200G

High Speed Adapter Forecast
2021 TAM (\$2.3B)



- Key Drivers of 25G+ Networking:

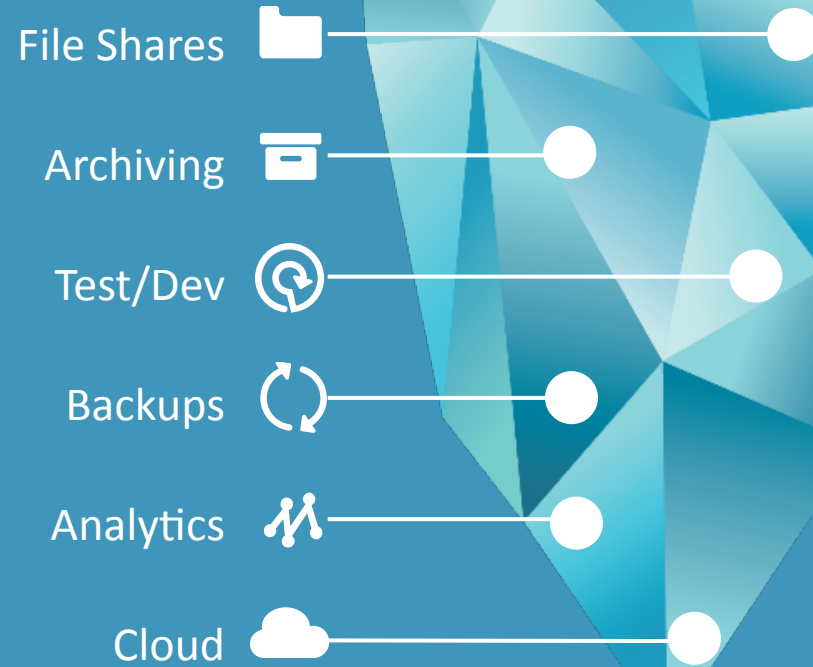
- Cloud, Big Data, Faster Storage, Virtualization, AI & Machine Learning

- 25G+ Adoption Cycle:

- #1: US Hyperscalers #2: BAT in China #3: Third wave of cloud, telco, and storage in US, ASIA, and Japan

Storage Landscape

STORAGE ICEBERG



PRIMARY STORAGE

- Traditional SAN
- Only 20% of capacity

SECONDARY STORAGE

- 80% of capacity
- Rapid growth
- Diverse data types
- Scale-out, Ethernet-based
- Tiered data

Storage Networking Background: Fibre Channel & Ethernet

1997

Feature	Fibre Channel	Ethernet
Bandwidth	1 G	100 M
Supports	Block	Block, file
Lossless	Yes	No
Cost	High \$\$\$\$	Medium \$\$
Cloud / HCI	No / No	No / No
Vendors	Several	Many
SDS / Scale-out	No / No	No / No

Yesterday: Storage Network = FC

- Fibre Channel offered best performance
- All interesting storage was tier-1 block
- No cloud or hyperconverged

2017



Feature	Fibre Channel	Ethernet
Bandwidth	8/16/32 G	10/25/40/100 G
Supports	Block	Block, file, object
Lossless	Yes	Yes
Cost	Medium \$\$	Low \$
Cloud / HCI	No / No	Yes / Yes
Vendors	2 / 2	Many / Many
SDS / Scale-out	Rare / No	Yes / Yes

Today: Both FC & Ethernet for storage networks

- FC option for Primary Block Storage
- Ethernet only option for all Primary & Secondary Storage (Block, Object, NAS, Cloud, Hyperconverged, Big Data)

Problem Solved With an Ethernet Storage Fabric

Ethernet Storage Fabric



*Everything a Traditional SAN Offers but ...
Faster, Smarter, & Less Expensive*

FAST

- Highest Bandwidth
- Lowest latency
- RDMA and storage offloads
- Native NVMe-oF Acceleration


SMART

- Automated Discovery & Provisioning
- Security & Isolation
- Monitoring, Management, & Visualization
- Storage-aware QoS


EFFICIENT

- Optimized Form Factors
- Just Works Out of the Box
- Flexible: Block, File, Object, HCI
- Affordable: SAN without the \$\$

Traditional vs. ESF—Where to Draw the Line?



Legacy DC – FC SAN




Modern DC – Ethernet Storage Fabric

Why Not Just Use Your Existing Switches?




Existing switches might not be designed for storage

- Are there enough available ports?
- Will the switches support new speeds – 25, 40, 50, 100GbE?
- Blocking? How many hops? How much latency?
- Are switches too big for storage or HCI clusters?



Need to Look To the Future

- Today's network purchases must last 3-5 years
- Fast enough for flash
- Flexible for future applications



How to Enable an Ethernet Storage Fabric

Ethernet Storage Fabric needs dedicated ESF switches



Performance



High Availability



Simple



Automated



Scalable



Cost Efficient



- ✓ **2 Switches in 1RU**
- ✓ **Storage/HCI port count**
- ✓ **Zero Packet Loss**
- ✓ **Low Latency**
- ✓ **RoCE optimized switches (NVMe-oF)**
- ✓ **NEO for Network automation/visibility**
- ✓ **Native SDK on a container**
- ✓ **Cost optimized**
- ✓ **NOS alternatives**

Open Ethernet 100/50/40/25/10G Switch Portfolio

SN2700 – 32x100/40GbE (64x50GbE)
Ideal ToR / Aggregation



SN2410 – 8x100GbE + 48x25GbE
25GbE → 100GbE ToR



SN2100 – 16x100GbE ports (64x25GbE)
Ideal storage/Database Switch
Highest 25GbE Density Per Rack




SN2010 – 4x100GbE + 18x25GbE
Ideal HCI ToR Switch




High performance, Multi-Tenant, Scalable Data Centers

Spectrum is Built for ESF


Congestion Management





Fairness



Avoidable Packet Loss



Guaranteed Low Latency



www.zeropacketloss.com

Ethernet Storage Fabric Must Support RoCE

- RoCE is RDMA over Ethernet
 - Bypass CPU
 - Increase efficiency
- RoCE has growing support
 - All Operating Systems
 - Many Applications
 - More Storage Arrays
- Common Use Cases
 - Storage (incl. NVMe over Fabrics)
 - Big Data
 - Video processing
 - Machine Learning / AI



Low Latency

DCB / PFC

Easy Configuration

Fast-response ECN

Guaranteed QoS

Strong Telemetry


Automated Mgmt.

High bandwidth


RoCE Support Done Right!



Application Blocked by the Switch



Application Blocked by the Switch



Summary: Choosing an ESF Switch



Better Performance



Right Form Factors



Easy Setup



Better Visibility



Tested End-2-End



Mellanox



Maximizing Efficiency of Software Video Streaming

Lower CPU
Utilization

80% Less
CORES*

Less Servers &
Network

40% Less*

Lower Power,
Cooling & Space

30% Less*

Highest
Throughput
& Quality

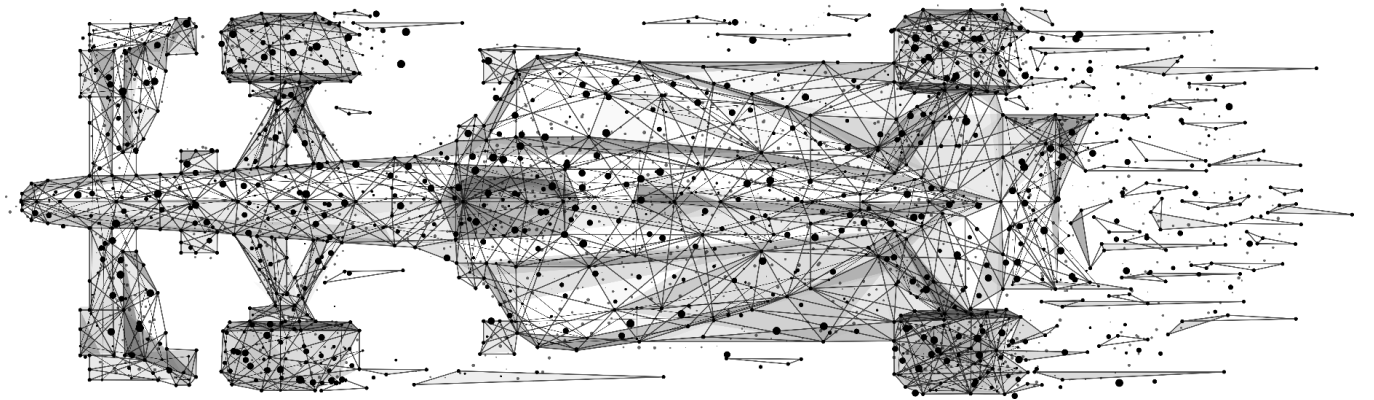
100 Gb/s
PER PORT

0% Packet
LOSS

2110-21
Compliance



Packet Paced in NIC hardware




* Comparing to Kernel based application performan

Ultra HD Resolution Demands Hardware Performance

- Raw uncompressed video bandwidth requirement exceeds 100Gb/s
 - Bandwidth for 1 video stream with 8K UHD high frame rate
- Packet pacing in software is not SMPTE 2110-21 compliant
 - Sub 1us Inter Packet Gap (IPG) is not achievable with software based solutions

Full HD	20bpp	25fps	IPG 11,000 nSec	Can be achieved with software packet pacing
4K UHD	20bpp	60fps	IPG 1,150 nSec	Cannot be achieved with Software packet pacing
4K UHD	24bpp	120fps	IPG 469 nSec	Cannot be achieved with Software packet pacing
8K UHD	36bpp	60fps	IPG 156 nSec	Cannot be achieved with Software packet pacing

Inter Packet Gap (IPG) Requirements for UHD



Bandwidth Requirements for UHD

Rivermax Key Features

Packet Pacing

- Leverages ConnectX-5 hardware based Packet Pacing
- SMPTE ST 2110-21 compliance at any bit rate
- No dependency on CPU Strength, OS interrupt level or Application



Kernel Bypass

- Reduced Kernel overhead with direct network adapter access
- Selective bypass – enables to select traffic bypasses and which flows to kernel
- Reduced latency
- Reduced CPU utilization
- Increased throughput




Packet Aggregation

- Application at Frame/Line(s) level
- Receive: fully assembled frame/lines(s) in memory
- Transmit: synchronously transmit packet paced full frames/lines (/chunks)

Packets vs **Frames**

Based on ConnectX-5 Technology



Thank You





NEXT-GEN CREATIVITY: BEYOND THE TECH

PRESENTED BY NICK ANDERSON

creative.space

powered by DigitalGlue 

MAKING 4K WITH HD TECH

STORAGE LIMITATIONS

- 1 & 10 GbE
- Non-realtime Image Sequences
- Compressed RAW for real-time


WORKSTATION LIMITATIONS

- Intermediate & Proxy Codecs
- Prores, Cineform, and DNx
- Optimized hardware acceleration



THE GREAT PLATEAU HOW FAR IS ENOUGH?

- 2K-8K Resolution (based on viewing distance)
- Physiological Color Gamut
- 21 Stops of Dynamic Range
7 stops visible at one time
- 12-bit



SO... WHAT CAN WE DO NOW?

- End-to-end Online Workflows
- Seamless Collaboration
- Artificial Intelligent Tools
- Masters as Deliverables
- ...and then it gets spooky

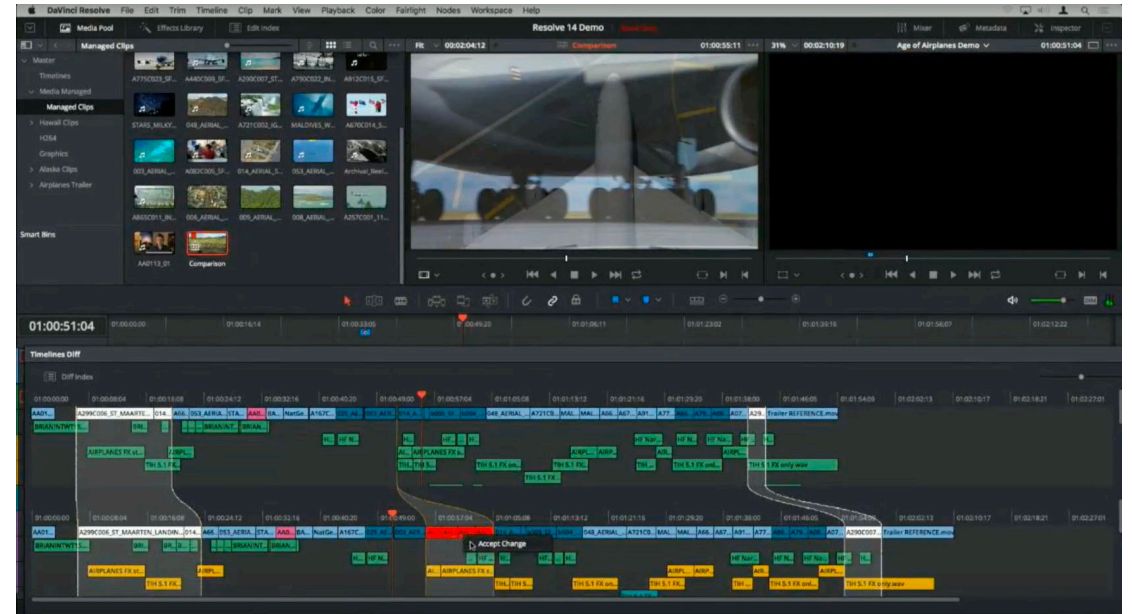
END-TO-END ONLINE

- No waiting for transcodes for source workflows
- Source to master quality instead of proxy (16-bit OpenEXR)
- More eyes looking at content to catch issues
- Non-destructive processing throughout the pipeline
- ACES Color Transforms



SEAMLESS COLLABORATION

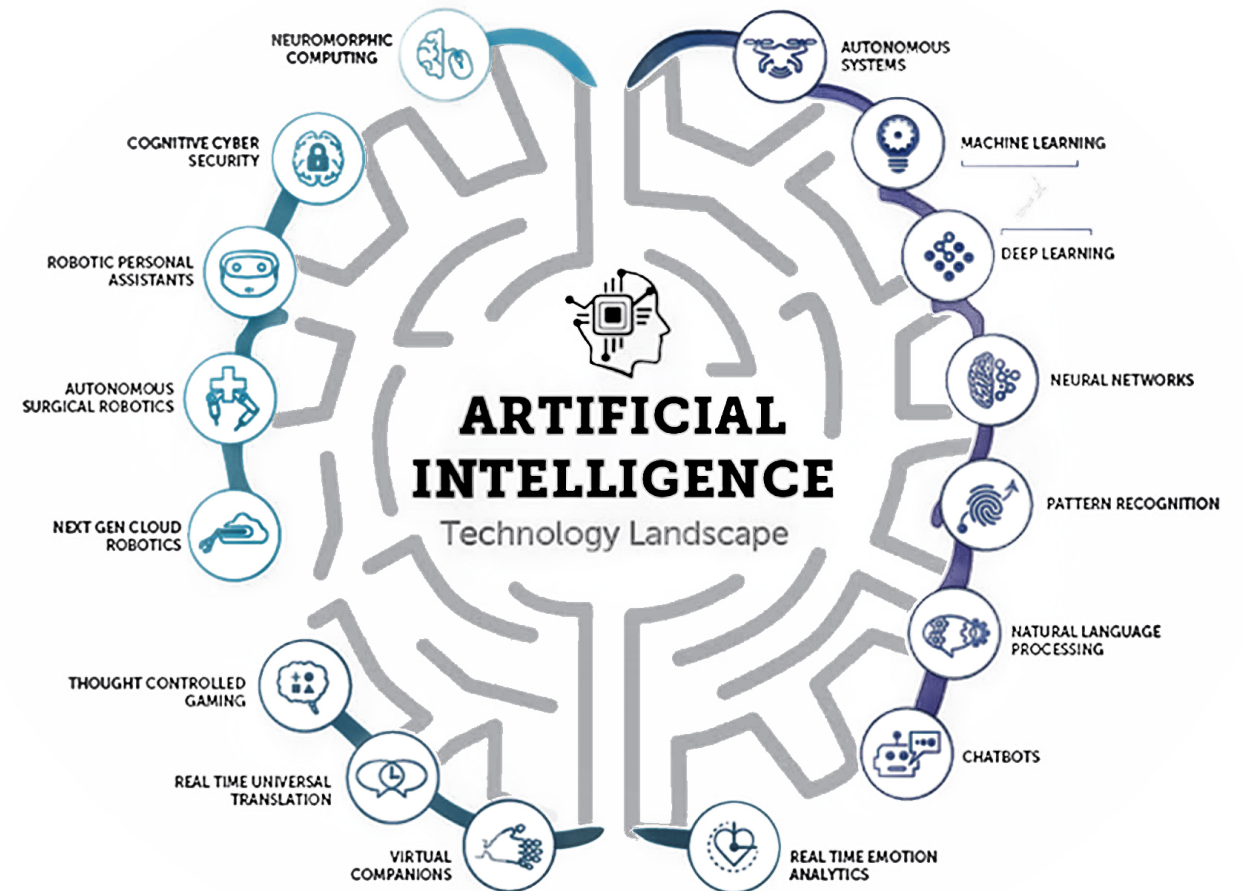
- NLEs with collaboration
 - ❖ DaVinci Resolve – Edit, Color, Audio, and VFX in one application
 - ❖ Adobe – Dynamic Linking with Premiere, After Effects, and Audition
 - ❖ AVID – Bin Sharing and Fusion Connect to DaVinci Resolve



- Entire team working in Parallel with the same assets
- No conforms, relinking, or project version tracking

ARTIFICIAL INTELLIGENCE FOR POST

- AI-friendly storage enables new workflows
 1. AI-based auto-tagging and transcription
 2. Creating new pixels and bits (Nvidia)
 3. Machine learning for automation
- AI designed to enable non-technical creatives




MASTERS AS DELIVERABLES

- High-bandwidth Fibre and 5G for content delivery
- IMF with encoding on the fly
- Metadata driven decode
 - Display transforms



THE COMING CREATIVE REVOLUTION

- Brain to computer interfaces
- P2P streaming
- Integrated AI creative tools
- Blockchain asset management
- Augmented, virtual, and mixed reality
- Femto-photography



THANK YOU

NICK ANDERSON
PRODUCT MANAGER, DIGITALGLUE
nick.anderson@digitalglue.com

creative.space[®]

powered by DigitalGlue 

powered by DigitalGlue 