

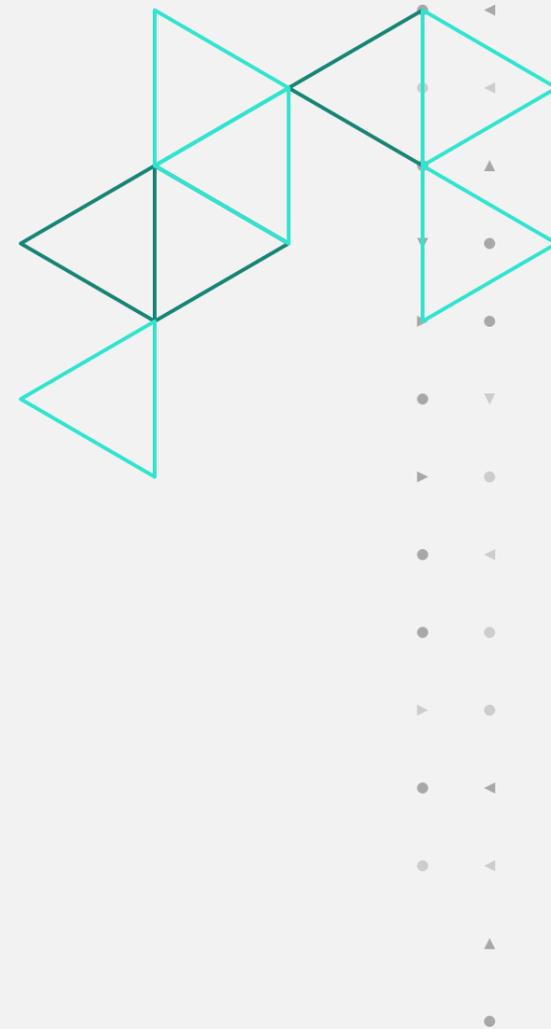


Microsoft Game Stack Live



DirectStorage for Windows

Andrew Yeung



Agenda

A DirectStorage for Windows update

The journey to eliminate load times

Spoiler: PC isn't there yet. Console is ahead.

DirectStorage: The next big step

Building a DirectStorage-ready engine

The journey to eliminate load times



HW fuels gaming innovation

Constant competition to build the best gaming HW

Feeds innovation in gaming experiences

- Resolution
- Frame rate
- Improved load times

Component

Processor speeds

Display resolutions

VRAM

IO speeds

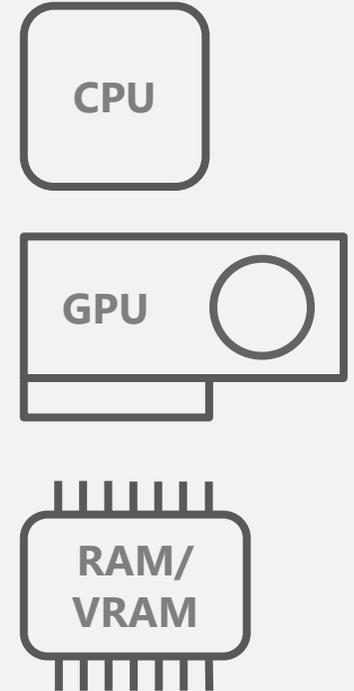
Evolution

KHz → MHz → GHz

720 → 1080 → 4k

KB → MB → GB

KB/s → MB/s → GB/s



GB/s still feels like MB/s

IO tech has improved so much, and yet...

- I only have a short time to play and a lot of it is just loading screens
- Every time I die, I have to wait another 20 seconds
- When I fast-travel, it's like I'm playing with early 2,000s textures again
- I sometimes feel like I bought an elevator simulator

Are we utilizing the HW effectively?



GB/s still feels like MB/s

TL;DR: IO HW has recently evolved beyond what SW can currently accommodate

HDD → NVMe SSD is a huge leap in IO capabilities

- 40x faster on Xbox Series X vs. Xbox One
- Same magnitude of innovation on PC

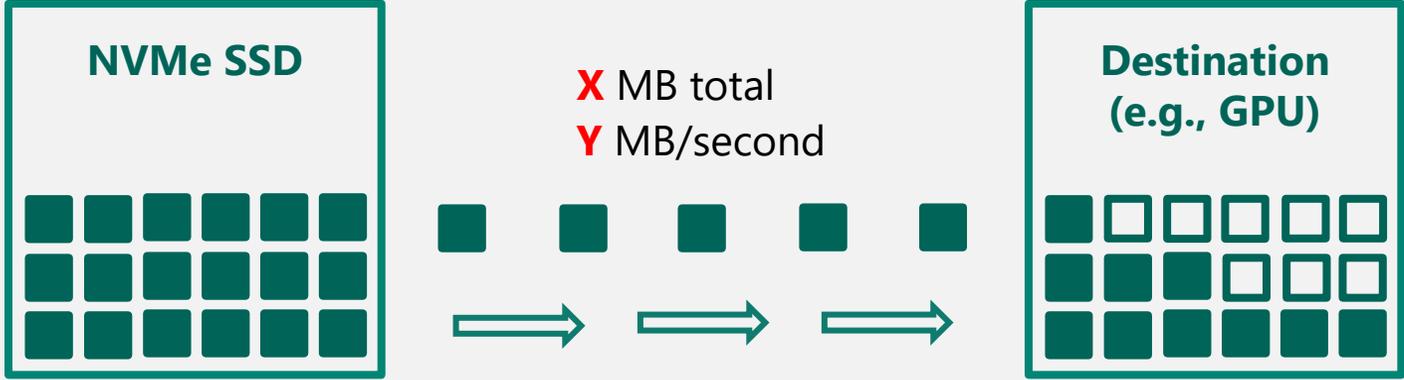
PC games run on NVMe SSD only see ~1.5–3x improvement in load times

Why?

- Not taking advantage of the full IO bandwidth (e.g., CPU bound)
- Too much data taking too long to load
- How do we get to the **<1 second load time** North Star?



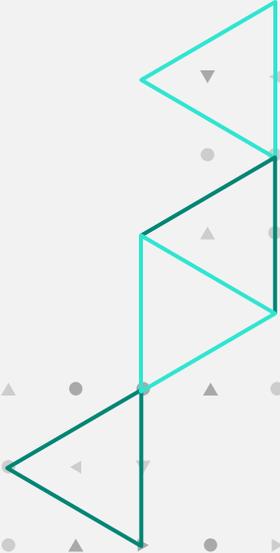
Getting next-gen speeds



Time to load data == X / Y

- Minimize X (**smarter** about what to load)
- Maximize Y (**faster** throughput)

Smarter ——— Sampler Feedback
Faster ——— DirectStorage



Maximizing HW potential on both Xbox and PC

Xbox Velocity Architecture lets you get the most out of your console's cutting edge HW

- All users with Xbox Series X/S consoles will benefit

PCs aren't all the same, but we still enable you to get the most out of the available HW

- New GPU's add features → gamers upgrade
- NVMe SSD prices drop → sudden but not instantaneous shift of users

One converged DirectX platform + cutting edge features = The best games on both console and PC

- DirectX 12 Ultimate
- DirectStorage
- --- Future innovations

DirectStorage: The next big step



Architecture

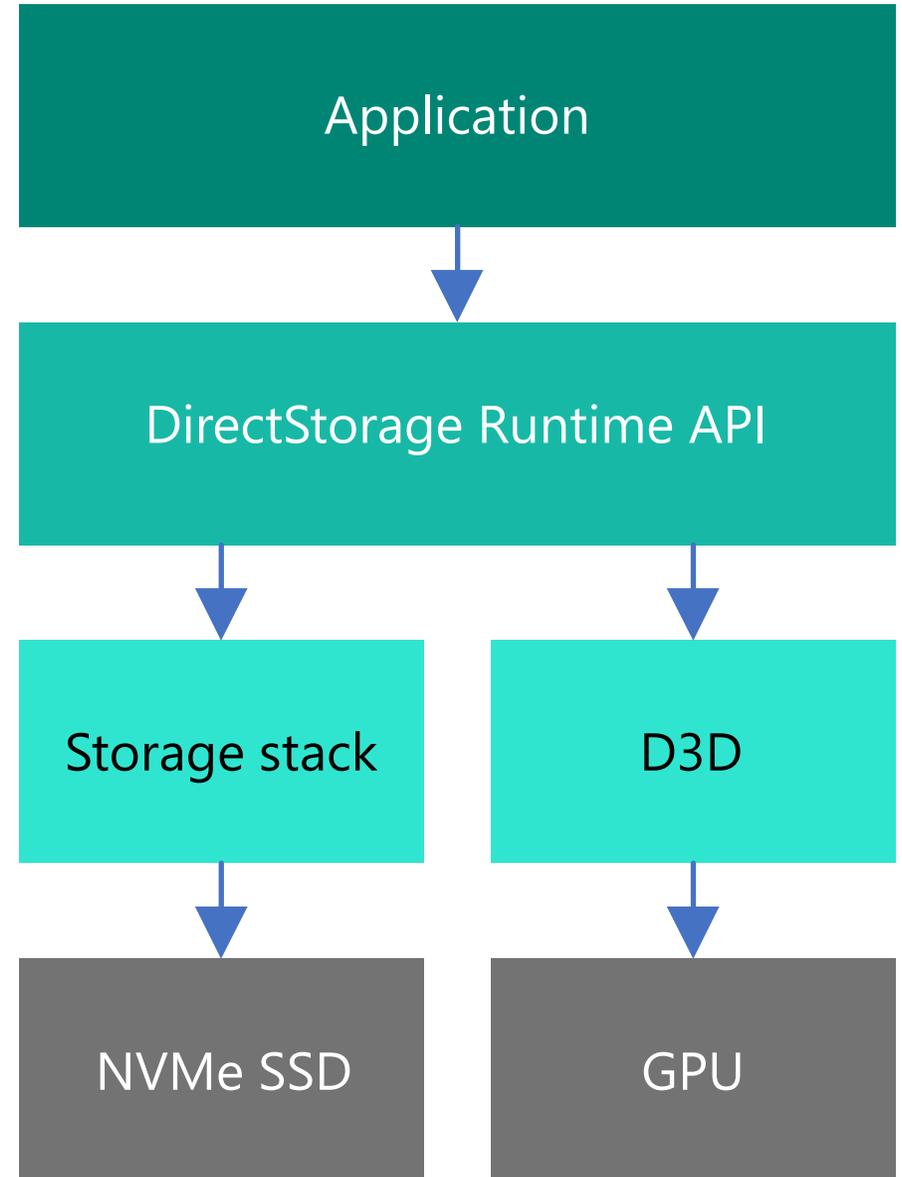
Goals and design principles same as Xbox

- API very similar with only minor differences
- Implementation details differ

Three primary components

1. DirectStorage Runtime + API
2. Asset decompression
 - a) Powered by D3D
3. Windows storage stack

These components work together to optimize end to end data flow

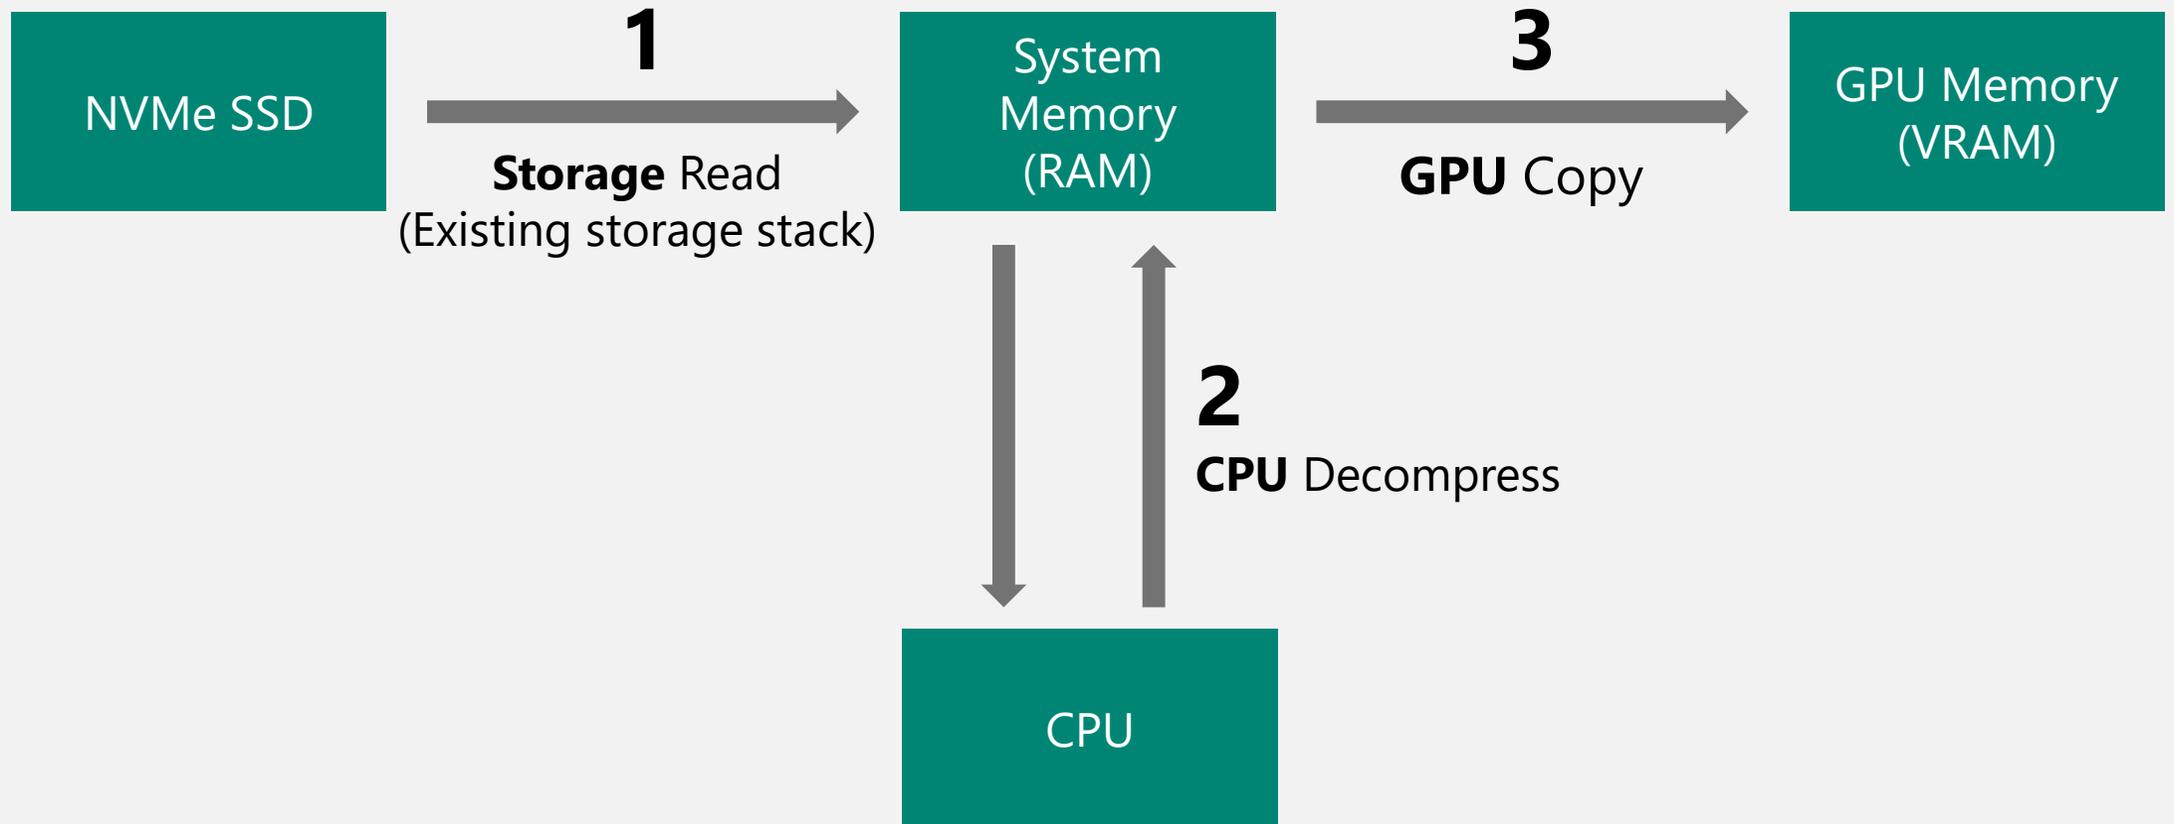


Scenarios

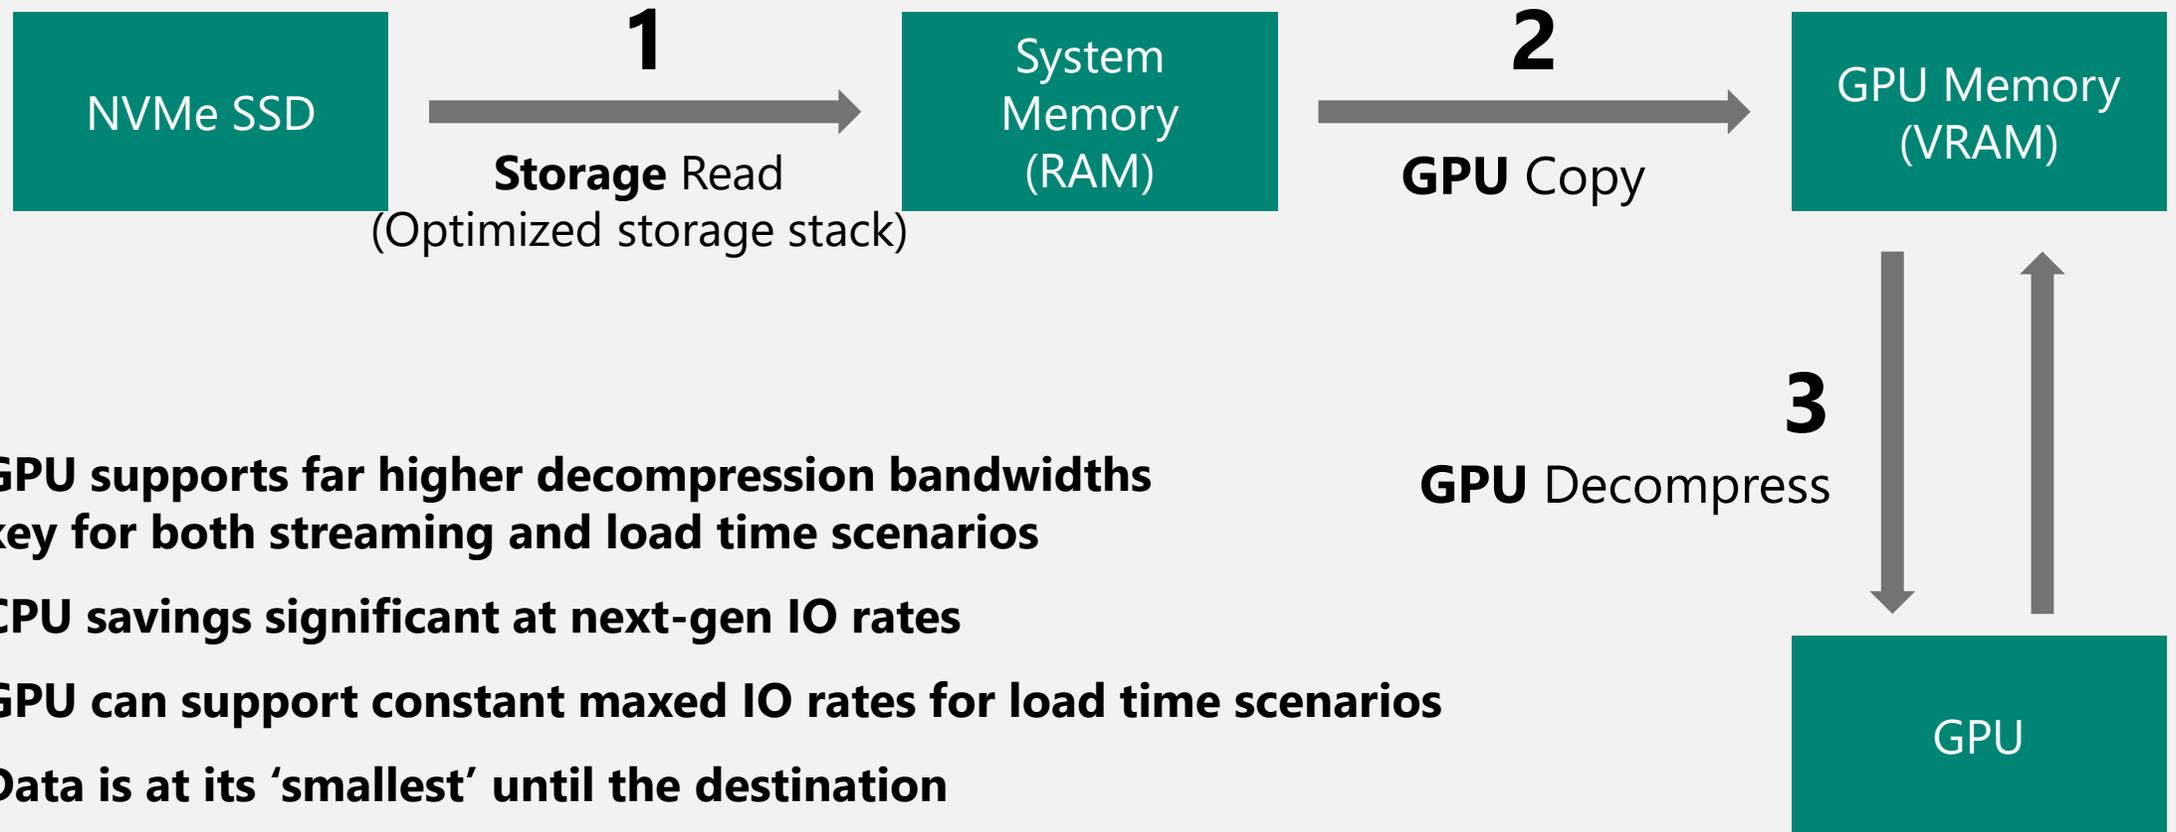
1. Copying data from disk for GPU usage
 - a) Geometry
 - b) Textures
2. Copying data from disk for CPU usage
 - a) Audio
3. Decompressing data for GPU usage
 - a) Necessary for assets like textures (e.g., DEFLATE compression on top of BC formats)



Flow of GPU assets (today)



Flow of GPU assets (with DirectStorage for Windows)



- **GPU supports far higher decompression bandwidths key for both streaming and load time scenarios**
- **CPU savings significant at next-gen IO rates**
- **GPU can support constant maxed IO rates for load time scenarios**
- **Data is at its 'smallest' until the destination**
- **Opens the door to other optimizations**

DirectStorage API

Newest addition to the **DirectX family** of APIs

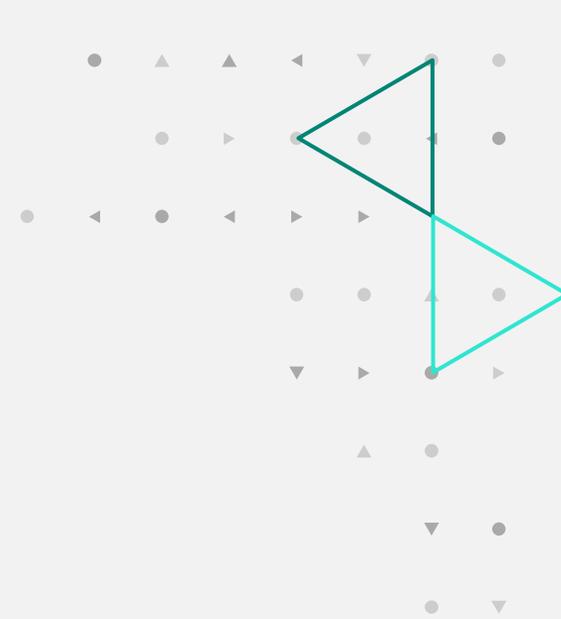
- Integrates directly with D3D12 maximizing the potential of both storage and GPU tech

New calling pattern paradigm that optimizes for 3 things:

- Keeping an NVMe SSD drive fed
- Keeping the GPU fed
- Reducing application overhead

Fundamentally, the future of high-performance scenarios is highly granular but very parallel

- These workloads like to have multiple requests upfront only notifying completion when necessary
- I.e. Batched workloads



DirectStorage API

Direct3D 12: Optimized model for leveraging GPU parallel behavior

DirectStorage: Introduces the batched model to storage read requests and bridges storage with GPU

- Maximizes NVMe SSD performance
- Implements most efficient path to connect data from disk to GPU

Existing IO APIs treat each request individually, with one completion notification per request

- Very inefficient for game type workloads with large # requests per second

DirectStorage APIs allow for developer control over creating batched requests and when to be notified of completions

- Significantly reduces application 'paperwork' to only what the app needs

Asset decompression

Gaming is a trailblazer always pushing tech to the limits

- The biggest, most detailed game worlds in 4K don't come for free (and what about 8k?)
- These need big asset packages that are compressed to fit on disk
- They also need to be decompressed at runtime as the game needs them

Specialized CPU based decompressors were created and have been good enough for now

- But at the higher rates of today, decompression overhead becoming too expensive

Gaming has created the need, but to make the next leap:

We need new innovation in compression tech

Asset decompression

Xbox Series X/S

- **HW accelerated decompression** as part of the Xbox Velocity Architecture to decompress assets

No such solution for PC

- How do you accommodate the NVMe SSD drive's high read speeds?
- Silicon takes time to build and won't be instantly available to everyone
- "Please wait a few years while we wait for HW to arrive" → **Unacceptable**

Instead....

GPU-friendly decompression

MSFT + GPU vendor collaboration to innovate and create a GPU friendly compression solution that works on **today's GPU HW**

- A new class of compression tech with new requirements

Sub-components under development

DirectCompute based decompressor

- Initial prototype decompression saturate gaming NVMe SSDs bandwidths
- Open door to further innovate in possible future silicon implementations

CPU decompressor for assets destined for system memory

Compressor

Windows storage stack

Optimizing Windows for high bandwidth + high IOPS DirectStorage scenarios

- Game usage patterns need high IOPS without sacrificing bandwidth
- Many small requests like texture tiles, not single GB size files

Achieved by optimizing path length & reducing CPU usage

- IOs bypass parts of the stack (e.g., file system and volume layers)
- Reduced user-mode \leftrightarrow kernel-mode transitions via batching/queuing
- Big focus on reducing per read CPU overhead

Pure Windows OS improvements. No developer action required

DirectStorage will automatically leverage the most optimized storage access path

Building a DirectStorage-ready engine



Sampler Feedback + DirectStorage

Increase in game memory needs faster than increase in VRAM on GPUs

Demands that games become **Smarter** and **Faster**

Memory multiplier techniques are a powerful tool

- Better memory efficiency through data selectiveness with **Sampler Feedback**
- Very fast asset turnover with **DirectStorage**

Makes it look like your GPU has x times more memory than what's physically available

Not just faster load times but also a path to bigger worlds and higher res detail only where you need it

Start building for next-gen now!

Begin rearchitecting your PC engines for next-gen IO patterns

- Start by integrating Sampler Feedback and DirectX12 Ultimate
- <https://devblogs.microsoft.com/directx/coming-to-directx-12-sampler-feedback-some-useful-once-hidden-data-unlocked/>
- https://www.youtube.com/watch?v=q7AG23s2_AM

Limited DirectStorage developer preview: Summer 2021

Reach out to us at AskWinDStorage@microsoft.com with your name, email, and studio name if you're interested in participating

Get updates on our blog: <https://devblogs.microsoft.com/directx/>

A huge thanks to our industry partners



We have a continued commitment to helping you get the most out of gaming HW in a **useful + standardized** way through **DirectX family APIs**

Helpful reading

Xbox Velocity Architecture

<https://news.xbox.com/en-us/2020/07/14/a-closer-look-at-xbox-velocity-architecture/>

DirectStorage is coming to PC

<https://devblogs.microsoft.com/directx/directstorage-is-coming-to-pc/>

Sampler feedback

<https://devblogs.microsoft.com/directx/coming-to-directx-12-sampler-feedback-some-useful-once-hidden-data-unlocked/>

https://www.youtube.com/watch?v=q7AG23s2_AM

