



### **Key Issues**

**Netflix has six months to build AVoD arm – mission impossible?  
AI analytics space smells streaming blood  
Still growing, still loss-making – another day at Disney+  
Mola TV trading video for blockchain in UK, discredits anti-piracy**

### **Analysis**

**Intel M&A spree continues, MaxLinear makes \$3.8 billion SSD leap**

### **OTT Video**

**Wynplay Community: Being ahead of the curve is a blessing and a curse**

### **Triple and Quad Play**

**RDK-B bolts on Airties' EasyMesh, as telcos consider common stack  
Dish Network defying financial odds, despite churning TV subs for fun**

### **Advertising**

**Full Throttle scales on coattails of desperate MVPDs  
Strategus guides brands through CTV ad maze, trashes walled gardens**

### **Sustainability**

**Tachyum tackles idle waste with reprogrammable Prodigy silicon  
Meta ruffles OpenAI feathers with carbon-crunching language models**

### **Worth Noting**

**OTT Video News, Deals, Launches and Products**

## Key Issues

Netflix has six months to build AVoD arm – mission impossible?

- **Despite its R&D reputation, Netflix has accepted it cannot go it alone in ads**
- **The Trade Desk is the first guiding hand for ad buying, maybe of many**
- **New tech, new teams, new trade-offs – internal disruption is inevitable**

Netflix will begin its contentious pivot to advertising as an alternative revenue stream from Q4 2022. In one of the greatest U-turns of the streaming era, the SVoD king will have to trade its fierce trendsetting status for a laggard wading into a marketplace where its vulnerabilities are waiting to be exposed by the competition.

According to the usually reliable New York Times, an executive note has been pinged to Netflix employees warning of this accelerated timeline for its ad-supported business before the end of this year. This is more than a slight acceleration – dramatically pulling forward plans that co-CEO Reed Hastings only recently said the company would “figure out over the next year or two.”

The message mentions that most of Netflix’s big US competitors have either planned to introduce ads or are running full steam ahead with ads to complement subscriptions, citing Apple, Disney, HBO Max, and Hulu, as per the NYT’s two anonymous inside sources.

“Yes, it’s fast and ambitious and it will require some trade-offs,” states the intercepted note.

By trade-offs, we assume Netflix execs are alluding to the \$70 billion that has been wiped off its market cap in the weeks since its first quarter results which exposed a loss of 200,000 subscribers. Trade-offs could also be referring to changes on a personal level, warning employees that an ad-supported Netflix could mean sweeping changes to how departments are run.

This could even go as far as lay-offs, to pay for the extensive new advertising team that it will require – from high-profile executives with extensive experience, to battle-hardened ad sales and marketing people, and not forgetting the best ad tech stack developers that money can buy. It will also need a dedicated data team for targeted advertising and an aggressive path to programmatic inventory.

Netflix may choose to align these teams and technologies closely with its existing SVoD operations, although there is certainly a case for keeping them separate, at least initially, to minimize disruption.

To execute all this in such a short period of time – less than six months – is almost mission impossible. Even if Netflix nails the technical processes of AVoD, we feel the rapid shift in strategy could risk inter-departmental rifts.

Netflix is renowned for its in-house technological developments, from its work on video compression and CDNs, to machine learning, recommendations, and analytics, but Netflix cannot go it alone in the unforgiving world of ads.

This is why Netflix is in discussions with California-based ad-buying platform The Trade Desk, according to these same NYT sources. Having former Netflix CFO David Wells on The Trade Desk board is the sweetener here, much to the disappointment of any other demand-side buying platform vendors hoping for a foot in the door at Netflix.

Of course, if the organization's foray into ad-supported streaming proves fruitful, there will be plenty more lucrative Netflix contracts up for grabs as it scales its buying power globally.

Besides, it is perfectly normal for a content powerhouse to recruit more than one demand-side platform for programmatic inventory, for example to provide transparency over ad order and age-specific advertising rules.

Netflix has a reliable supplier in The Trade Desk, which processes some 800 billion ad impressions on a daily basis, working with pretty much every ad-supported network you can think of. One recent example is becoming the first ad exchange to programmatically integrate Fox's VoD inventory.

The sensible option will be to introduce the new ad-supported tier at a price point of around \$10 a month, in line with HBO Max's ad-supported offering. The more aggressive option, which we hope Netflix takes, would be to undercut the competition in a big way to see how the market responds.

By starting low, this would allow Netflix to incrementally increase pricing once momentum picks up, although that evokes memories of how the live streaming skinny bundles – like Sling TV and YouTube TV – went to market priced way too low, and caused consumer uproar when prices were raised to level required to turn a profit.

Who knows, Netflix may even throw a curve ball by offering its ad-supported option totally free (albeit unlikely).

**We feel the rapid shift could risk inter-departmental rifts**

**Netflix can afford to be aggressive in ads – but will it take a risk?**

We shouldn't forget that Netflix is still the out-and-out market leader in SVoD. It can afford to be aggressive and throw its weight around. Advertisers will be flocking to the platform and paying a premium, soon forgetting that Hastings and co. were ever vocal critics of advertising.

The NYT's note that keeps on giving has also fleshed out additional details on Netflix's password sharing monetization plans, with higher charges being introduced for active password sharers also around Q4 this year.

NYT sources weren't in a position to provide guidance on by how much prices might be hiked, but it's likely the \$15.99 monthly standard tier could reach \$20 for password sharers. This is guaranteed to cause cross-state family arguments and trigger an initial spate of cancellations, but it might also normalize subscription bill-splitting akin to what friends might do in a restaurant.

Plenty of users do that already, considering that Standard gets your two simultaneous screen and Premium (\$19.99) gets you four simultaneous streams. Against this background, the password sharing crackdown has been a difficult pill to swallow, with Netflix essentially charging subscribers more for something many will feel they are already entitled to with current subscriptions – multiple screens regardless of postcode.

These plans are also an acceleration from guidance provided during Netflix's earnings call just weeks ago, when execs said password sharing monetization mechanisms would take "a year or so of iterating."

From the cherrypicked snippets of this small but powerful note delivered into NYT hands, we feel important details have been conveniently omitted by sources. The part on "trade-offs" is a particular eyebrow-raiser, with connotations that seem world's away from those paraded in Reed Hastings' book, *No Rules Rules: Netflix and The Culture of Reinvention*.

The co-authored business book published in 2020 is exactly what it says on the tin – a tale of workforce liberation masquerading as a vanity project.

The argument is that Netflix was propelled to the award-winning big leagues by allowing people to work when they want, how they want, take holiday when they want, and be paid whatever they want, with the caveat that they deliver absolutely perfect results on time.

Only because of this culture has Netflix been able to reinvent it so successfully, or so the story goes, although we feel that segueing from SVoD to AVoD represents a bigger challenge for Netflix than its fabled pivot from DVDs to the world wide web.

## Key Issues

### AI analytics space smells streaming blood

Video analytics vendors are using recent shifts in the streaming landscape to their advantage, twisting a narrative of 'peak streaming' whereby only OTT content providers with the largest datasets and strongest AI-based platforms will survive.

While the industry should do its best to ignore such sensationalist headlines related to Netflix's subscriber blip and the implosion of CNN+, it is true these two incidents have contributed to a new analytics landgrab as service providers scramble to harness everything there is to know about user behaviors, out of fear they might be next.

Symphony MediaAI, for instance, is spinning up a marketing drive based on these high-profile turning points for global streaming. The US analytics firm's CTO, S.V. Vasudaven, has undergone a drastic mindset change as a direct result of the recent weaknesses of Netflix and CNN+.

Speaking on a webinar panel on AI and analytics, hosted by IBC365, Vasudaven claims his message would have been very different if the event was taking place a month earlier. Now, he is waxing lyrical about a surge in competition for consumer attention, particularly in acquiring and maintaining that attention.

On one hand, this may be an opportunistic PR stunt on Symphony MediaAI's behalf. On the other, it might be a genuine reflection of customer demand that we should be taking seriously.

What's clear is that Symphony MediaAI believes the next frontier in video analytics lies in delivering AI-based tools to data teams to predict and execute on game-changing business decisions well ahead of the competition.

Holistically, Symphony MediaAI's software platform helps content provider partners manage, track, predict and grow content monetization. The manage and track parts of the platform are what Vasudaven describes as the fundamental nuts and bolts of accounting – such as tracking where revenue is coming from, which distributors are providing it, and who's performing best.

**One client was left red-faced after uncovering a cluster of subs on perpetual free trials**

Once you are able to capture the financial nuts and bolts, you are in a position to be able to predict with what he calls ‘classic analytics,’ which relates to questions such as why are things trending? Or which distributors are working better for me than others?

This classic approach to analytics can tell content providers where they are heading, but the difference with AI-based analytics is that you can discover where you *should* go. AI in this context can answer questions like which piece of content should I license? Which content will perform well with which audience? Are there signals I can pick up on to learn where strategically I should head?

AI is good for helping with tasks such as classification, predicting, detecting outliers, and decisioning. This issue is that there’s so much data that has come into the situation that you need machine learning to help out.

Vasudaven cited a use case of using AI/ML to mitigate churn – working with an unnamed SVoD provider to understand and predict when subscribers will churn.

In this case, the client provided Symphony MediaAI with a wealth of signals, which it ran through ML algorithms and found behavioral characteristics that indicated tendency to churn, including binge watching certain titles before churning.

When a user starts bingeing one of these titles, this should then be a trigger to call your best churn-prevention mechanisms into action.

One data grouping even found a cluster of subscribers that had been on perpetual free trials due to a software bug, which was an embarrassingly obvious missed revenue opportunity that Symphony MediaAI helped to monetize. No wonder this customer doesn’t want to be publicly named.

This ties into how the platform is becoming more proactive in nature, such as by identifying subsets of viewers with a high probability of churn and offering a number of actions to circumvent this.

“In an unsupervised ML environment, you can capture the shape of data,” Vasudaven explains. “But ML knows nothing about video, just data; grouping data together which is similar and allowing us to go in and inspect these groupings. This surfaces what is referred to in the industry as unknown unknowns.”

Symphony MediaAI serves both subscription-based and ad-supported clients with very different needs. Its tools light up insights into how consumers are interacting with content, but with a real focus on where the money is coming from, reducing expenses around customer acquisition, and focusing on maximizing existing customers, thanks to its parent company's roots in financial services.

What impresses us about Symphony MediaAI is how it melds two data reservoirs that traditionally are accessed separately. On one side, it works with content providers on accounting, where subscriber churn burns a gaping hole on the balance sheet, most of the time. But by only analyzing accounting, you miss insights into why consumers have churned.

Symphony MediaAI is still in the early stages of marrying money tracking with behavioral data, to answer questions such as what is a consumer actually doing? What are they watching? How much are they watching?

It's only when you bring these two datasets together that you can see what decisions to make on content procurement and what effect it has on viewership. Did I make the right choices? How can I close the loop on the decision process, to see the cause and effect relationship between business choices and consumer response?

For Vasudaven, one of his favorite aspects of AI is that once you have the history, Symphony MediaAI can provide tools to content executives to make growth predictions. He has a prediction of his own, foreseeing increasing demand over the next several months for AI and ML capabilities in addition to the company's base product, to surface insights that are humanly impossible to visualize without computational horsepower.

"It's going to benefit consumers to access content they're interested in and will benefit content providers to make the right investments to maximize enjoyment of subscriber populations," voiced Vasudaven.

All this talk of AI and ML wouldn't be complete without an obligatory contingency comment from Vasudaven on how these tools are designed to support data science teams, not represent the threat of replacement.

Symphony MediaAI has seemingly come from nowhere, really as two companies in one that has burst onto the media scene from the financial services side where it has operated for around 30 years. As the name suggests, the focus is on media, but it has sister companies focusing on areas such as healthcare AI and retail AI.

**"ML knows  
nothing about  
video, just data"**

## Key Issues

### Still growing, still loss-making – another day at Disney+

Disney+ added 7.9 million new subscribers in Q1 2022 (Disney's second fiscal quarter) to total 137.7 million users globally, including the Disney+ Hotstar platform in India.

With Hulu now on 45.6 million subscribers and ESPN+ on 22.3 million, this brings Disney's streaming empire to 205.6 million users across its global direct to consumer streaming properties.

Yes, Disney is closing the gap on Netflix (221.6 million subs), and yes, Disney will take top spot soon enough, probably by the end of this year – yet this is a high-level, napkin-math view, as shown in the graph below.

We would love to get more granular than this, looking at how many individual subscribers Disney actually has, and how many of these are signed up for multiple video products, but that is currently impossible.

While we know for certain that one Netflix subscription is counted as one subscriber, even if the same account is being used across multiple households, Disney does not adhere to the same accounting principles.

If a household subscribes to a bundle of Disney+, Hulu, and ESPN+, this is counted as three separate subscribers, which is both misleading and totally understandable.

We would much rather Disney break its services out separately than only provide one global, bundled number, which it might decide to do one day.

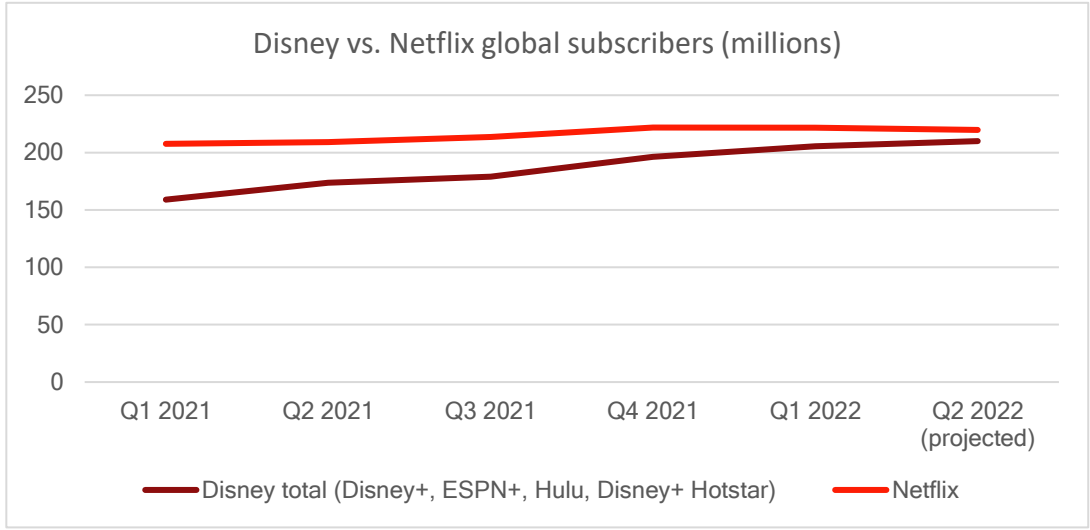
Of course, Disney is not immune to password sharing either, but the company isn't being nearly as transparent as Netflix in terms of sharing the scale of the missed opportunity, which the current SVoD leader estimates as approximately 200 million subscribers.

From a year on year perspective, ESPN+ has performed strongest among Disney's D2C businesses, growing its user base by 62% between Q1 2021 and Q1 2022, virtually double the 33% annual growth seen at Disney+ and dwarfing the 10% annual increase in Hulu subs.



More good news is that ARPU has increased across the board for all D2C streaming services. The biggest ARPU gains came on the Hulu Live TV + SVoD services, increasing from \$81.83 to \$88.77, while the rise at Disney+ Hotstar – although tiny – is the most significant by far, spiking from \$0.49 to \$0.76.

This has had a positive impact on total group revenues – rising 23% to \$19.2 billion for the quarter.



Within the Disney Media and Entertainment Distribution segment, where revenues increased 9% to \$13.6 billion, D2C sales accounted for \$4.9 billion, a solid increase of 23%. However, Media and Entertainment income fell 32% to \$1.9 billion – driven by a 205% increase in operating loss to imprint an ominous \$887 million hole on the latest quarterly income statement.

Disney blames the increase in operating loss on higher losses at Disney+ and ESPN+, as well as lower operating income at Hulu. In other words, all its D2C streaming businesses are still loss making, and will be for a very long time, while Netflix is making tons of money.

### Key Issues

#### Mola TV trading video for blockchain in UK, discredits anti-piracy

Faultline’s call with Mola TV this week was full of surprises. The Indonesian OTT video and pay TV outlet is best known for snatching soccer viewing rights away from the status quo, in a similar disruptive vein to DAZN, although - unbeknown to many - Mola TV also has a budding blockchain-based events business and plans to enter the gaming space.

Mirwan Suwarso, Mola TV’s official representative, cut a refreshingly honest character from the outset – illuminating us to the company’s content roadmap while throwing considerable shade over the state of the anti-piracy market.

Pre-briefing, we head over to the Mola TV website for a nose around, only to be met with the message, 'Our services are under improvement at the moment' – which seemed like an apt place to start.

"We've pretty much given up on the UK website," Suwarso admits, without hesitation. "How do you compete with something like BBC iPlayer? It's fantastic and it's free."

Mola TV announced its operational expansion into the UK and Italy in October 2021, but things have not gone quite to plan from a direct to consumer streaming perspective.

This website apathy should not be mistaken for a total exit from the UK market, however, as Mola TV has identified the UK as the perfect place for its blockchain-based activities.

While Mola TV's decentralized operations are currently focused on events, the company is exploring how blockchain can be complementary to its core video streaming business.

**"How do you compete with something like BBC iPlayer?"**

The idea of simplifying the delivery chain while achieving scalability via blockchain is something that still sounds too good to be true, yet NFTs have flown out of the blocks post-Covid, as a way of commodifying digital assets such as video and music files, that are uploaded to a blockchain where copies are recorded on the digital ledger.

These NFTs can be purchased with cryptocurrency, without the purchaser owning the original file, so artists retain copyright. Unlike cryptocurrencies, however, the totally unique nature of NFTs means they are not interchangeable with any other asset.

By making blockchain-based decentralized delivery networks more widely available to content creators and media partners, organizations like Mola TV will be able to reduce overheads in areas such as transcoding and delivery fees, but also create new monetization opportunities in areas of advertising, ticketing, and merchandise.

Skeptics are still in abundance in the TV ecosystem, Faultline being one.

Yet it's for this reason Suwarso can claim confidently that Mola TV has no proper competitors, holistically speaking, because a rival like DAZN doesn't have a blockchain-based events business to fall back on, when streaming live sports fails to make any money.

Suwarso warns this is a long game. The blockchain-based systems being developed in-house at Mola TV should eventually be built into the video platform by Q4 2023. A lot of experimentation is needed in the meantime before this long-term vision becomes a blockchain reality, if at all, and mistakes will be made along the way.

A more pressing matter for Faultline's typical audience, however, concerns Suwarso's outright dismissal for spending big on anti-piracy.

This was particularly pertinent given the huge levels of piracy seen in Mola TV's motherland or Indonesia, as well as the fact Mola TV has contracted Verimatrix for multi-DRM, a deal which Suwarso sees as purely a contractual requirement of studios, rather than an organically proactive decision from the Mola TV team based on promises that deploying multi-DRM will generate genuine revenue savings.

Suwarso is far too thick-skinned from his experience in the futile fight against content pirates in Indonesia, Malaysia and Singapore.

Indonesia is among the world's fastest growing internet markets, as well as being one of the world's largest non-paying markets. "We will find a way not to pay," said Suwarso.

Despite this, he reveals that Mola TV spends around \$9 million a year on combatting piracy – but not on technology, on deterrence.

"We laugh at the English Premier League (EPL) when it claims to help with piracy. It's no thanks to the rights holders we have to do everything ourselves. UFC does a way better job of anti-piracy than the EPL," continues Suwarso's tirade on anti-piracy.

He doesn't let up, telling us that while the Apple Store is basically void of piracy, you might as well forget about the Google Play Store.

As well as Verimatrix, Mola TV works with Viaccess-Orca, but we understand this deal is for video player and ad insertion capabilities, although some security is likely to be baked in given the French firm's content protection roots.

Akamai is Mola TV's preferred CDN, while Ateme also supplies some content delivery components via its Anevia CDN assets, along with a full suite of Titan encoders and packagers for live 4K UHD streams of EPL soccer matches.

**"We laugh at the EPL when it claims to help with piracy"**

Harmonic was first in the door for video compression at Mola TV, serving up its VOS360 cloud streaming infrastructure for UHD HDR streaming of live sports. The SaaS technology simplifies stages of the media processing and delivery chain, and Harmonic's EyeQ AI-based encoding technology improves streaming quality for Mola TV's streaming service.

We learned that Mola TV has an in-house recommendations engine and handles subscription services internally too, putting these components in pole position for its future blockchain video integration endeavors.

While the emphasis of our conversation is on streaming, as is Mola TV's marketing, we should note that legacy DTH is big business for Mola TV. It currently has over 10 million linear TV viewers in Indonesia through partnerships with operators, and is on the cusp of expanding into Malaysia and Singapore through telco partnership. The D2C OTT business, meanwhile, is currently on around 1 million active users.

After three years in the OTT market, you can sense subtle tones of disappointment about where Mola TV is today in terms of global expansion. There is also a more prevailing sense of optimism that Mola TV has identified the future, even if it is way too early to tell.

**“We don't fight for rights, we create our own rights”**

Founded in late 2018, Mola TV picked up EPL broadcasting rights in Indonesia as well as a treasure trove of local soccer content, with financial backing from parent company Polytron (a subsidiary of Indonesian manufacturing giant Djarum Group), before the platform launched officially in Q3 2019.

Arriving to market so soon before the pandemic was both a blessing and a curse, as Mola TV was sent back to the drawing board with all live sports placed on lockdown. The on-demand catalog was bulked up with kids content and the platform also branched into music, contributing to a huge ranging spectrum of audiences.

We mentioned DAZN earlier as a key competitor, but Suwarso doesn't see it this way. “We have sold some content to DAZN – we don't fight for rights, we create our own rights,” he refutes.

By this, he means Mola TV has ambitions of becoming a largely self-sufficient original production outfit. It has full production capabilities and studio assets in place, investing strongly in local soccer content as well as MMA, which Suwarso believes is the world's fastest growing sport.

With MMA, Mola TV is just dipping its toes into the Italian marketplace. This immediately evokes nightmares of DAZN's Serie A streaming calamity. Suwarso claims Mola TV has the streaming infrastructure in place to avoid a repeat banana skin situation, pointing to having up to 8 million concurrent streams for the Indonesian national soccer team's matches, which it took away from FTA, to much initial outrage.

He assures us that Indonesia's internet infrastructure is far worse than Italy's, although mobile-first viewing is more prevalent in Indonesia, while the big screen is the first port of call for live sports in Italy.

Seeing some traction in streamed events with interactive features such as fan Q&A sessions, this segued Mola TV into distributed ledgers and non-fungible tokens (NFTs), having conducted its first fully NFT ticketed event in London in March, which it sold out.

This is part of the new Mola Chill Club, an NFT-based membership where subscribers can access exclusive NFT tickets for concerts with VIP perks up for grabs, such as backstage passes and flight tickets.

Faultline's skepticism about the credibility of NFTs as a standalone business within the video streaming landscape were met with comments about some of Mola TV's recent NFTs tripling in value. It remains to be seen how this will be factored into Mola TV's live and on-demand OTT video offering, possibly as a rewards-based token scheme. Although the fact that the UK streaming website is being phased out in place of a blockchain-based events business has made us think twice.

## **Analysis**

### **Intel M&A spree continues, MaxLinear makes \$3.8 billion SSD leap**

Two silicon acquisitions have highlighted the continued tumult in the semiconductor industry, this week, at opposite ends of the scale. MaxLinear has made a CPE play, buying Silicon Motion and its SSD expertise for some \$3.8 billion, while Intel has bought a mobile-focused GPU specialist to help it brace against AMD and Nvidia in both the data center and new video experiences.

For MaxLinear, this is about being able to tick another box when it replies to equipment RFPs. To some extent, it is not an exciting deal. Meanwhile, Intel's recent run hints at two possible readings – depending on how much faith you place in the new regime.

**As data centers power the video process from beginning to end, developments are of interest**

Either it is continuing to flail, throwing cash at its myriad problems as it continues to lose market share to AMD's resurgence in CPUs, Nvidia's continued dominance in GPU, and now the rise of ARM-based chips in the data center; or it has finally refocused the tens-of-billions-of-dollars that it has spent in the past decade trying to pivot from its Xeon-based dominance and then shore-up data center incursions from new architectures.

As so much web traffic consists of video, and these data centers power this entire process from beginning to end, developments in the landscape are always of interest. It has looked, for a long time now, that Intel was on the verge of a major upset, but the industry has been slow to embrace new silicon architectures.

Intel is on something of an acquisition spree, as its new CEO Pat Gelsinger starts righting the perceived wrongs of his predecessors. In February, Tower Semiconductor was bought for \$5.4 billion, then March saw AI-based workload optimizer Granulate bought for \$650 million, and now May sees Siru Innovation's GPU expertise bought for an undisclosed fee.

Siru Innovations is a Finnish firm that can trace its heritage back to the 1990s – inside Bitboys, a desktop GPU frontrunner that eventually moved into mobile GPUs in the early 2000s.

ATI acquired it for \$44.5 million in 2006, which in turn was purchased by AMD two months later. In 2009, AMD sold its mobile assets to Qualcomm for \$65 million, and so the Bitboys founders left Qualcomm in 2011 to found Siru.

As you can see, the founders have been around the block, but now find themselves inside a giant that has inexplicably struggled to create a proper desktop GPU platform for years. The Intel Arc GPU portfolio is expected to be revealed shortly, which might finally allow it to challenge Nvidia and AMD, but like many of Intel's AI-focused silicon projects, there is a pretty decent chance that Arc will be a complete dud.

Siru's arrival is too late to alter the fate of the first generation of Arc products, but Intel did poach AMD's GPU head Raja Koduri in 2017, who has since been heading up the GPU project – and who will have worked with many of the Siru staff previously.

The cynical view is questioning what Siru brings to the table, as it has been mobile-focused for so long, but the optimistic one is that Intel might finally break the GPU duopoly that has plagued the x86 industry for the past two decades.

For Faultline's interest, the Arc product is arriving in a backwards fashion. The current Intel GPU roadmap targets laptops in Q1, desktops in Q2, and then workstations in Q3. Samsung is already selling some laptops featuring the discrete Arc chips, and the first desktop graphics cards are going to be sold to OEMs and system builders first, before being released to the public.

The workstation and server cards are what could prove transformative to video workloads, and this mobile GPU interest should see Intel expand into metaverse-y AR and VR devices.

Intel has spent billions on AI-focused R&D and acquisitions, which should also be brought to bear in the enterprise markets. One such example is Habana Labs, which was acquired for \$2 billion to provide silicon counters to Nvidia's GPU-based capabilities.

Two second-generation products from Habana have been unveiled this week, with the Gaudi2 and the Greco. Designed for data center deep learning and inference jobs, Intel has proudly proclaimed that they outperform Nvidia's A100 platform.

However, the A100 is getting on for two years old, and so is due a refresh. The H100 is expected in Q3, which is going to undermine Intel's current claims.

Silicon arms races are expected, of course, and the different software and architectures used in each workload influence performance far more than the raw hardware and firmware capabilities.

Intel's aforementioned Tower Semiconductor deal saw it acquire a "leading foundry for analog semiconductor solutions," which has focused on RF, power, and electronic design automation. This analog vein brings us to MaxLinear's \$3.8 billion acquisition of Silicon Motion, which will see the analog purveyor add Silicon Motion's range of NAND flash controllers for solid state storage to its lineup.

MaxLinear supplies many of the broadband and networking equipment vendors, as well as a whole pile of CPE, and so the move into memory controllers is not as odd as it first looks. It is always surprising to see how little attention is paid to all the auxiliary components packed alongside a CPU or GPU on a device's PCB, and that supporting-cast world is the one that MaxLinear lives wholly inside.

MaxLinear reckons the combined entity will have an enterprise value of some \$8 billion, with annual revenues of over \$2 billion, and a total market opportunity of some \$15 billion. Both

**We have a sneaking suspicion that MaxLinear will power Comcast's Xfinity DOCSIS 4.0 gateway**

Qualcomm and MediaTek were reportedly sniffing around Silicon Motion, and MaxLinear appears to have had to pay a premium to beat them to the punch – with the offer set at around a 48% premium over the closing share price.

With annual revenues of \$922.1 million, and operating margins in the 27% range, Silicon Motion should certainly not be a burden. The listed savings from combining the two are expected to hit \$100 million annually, which sweetens the deal, but again, this looks like a lot of cash to pay. The Wall Street crowd seem to agree, with Silicon Motion's shares jumping 17% after the announcement, but MaxLinear's falling 22% - the largest drop since March 2020.

Of course, the stock markets as a whole are taking a beating currently. MaxLinear is betting that it can save cash by in-housing its supply of memory controllers. Solid state flash storage is now commonplace among networking equipment and consumer electronics, and so MaxLinear gets to offer another slice of the pie in its negotiations with OEMs.

Currently, Faultline readers are most likely to recognize MaxLinear from its work supplying WiFi chips to CPE vendors, with CommScope, Compal, and Skyworth among its known customers.

As for operators, we have a sneaking suspicion that MaxLinear has a chance of being deployed in Comcast's Xfinity DOCSIS 4.0 gateways, with Broadcom another contender. We know that Quantenna's QSR10G WiFi chipsets power the Xfinity DOCSIS 3.1 gateway, while the accompanying Plume WiFi pod extenders used by Comcast run Qualcomm 2x2 WiFi chips.

Back in April 2020, MaxLinear snaffled Intel's home gateway business for some \$150 million, which looked like a bargain that had the potential to double its quarterly revenues. Since then, DOCSIS 4.0 work has generated buzz, as well as a possible resurgence in MoCA and its continued threat in the G.hn sector – after it acquired Marvel's G.hn portfolio for a paltry \$21 million back in 2017.

## OTT Video

### Wyplay Community: Being ahead of the curve is a blessing and a curse

A calendar mix-up at NAB 2022 left Faultline with unfished business at Wyplay – as our five-minute cameo raised more questions than answers about the French firm's new blockchain-based façade.



A follow-up call with co-founder and CMO Dominique Feral this week set our mind at ease about why Wyplay is making this potentially perilous pivot into cryptocurrency, and what actually constitutes the platform, called Community.

However, our confidence in the business model – whether Wyplay can actually make any money from minting non-fungible tokens (NFTs) and peddling decentralized infrastructure – is more skeptical than ever.

Wyplay Community has three pillars – Wyplay P2P, Wyplay Distributed DRM, and Wyplay Blockchain. Just by looking at these, before diving into the technical details and assessing the business case, you can see the problems.

No one makes money from DRM anymore, while creating decentralized networks, where users share compute resources in return for rewards, is certainly not a cash cow for a company like Wyplay, nor its traditional operator customer base.

Decentralizing infrastructure can reduce overheads from transcoding and delivery networks, although this is minimal in the grand scheme of things. Feral admits that delivery makes up less than 10% of the total cost of delivery content, so reducing delivery costs is not a core selling point of Wyplay Community.

So, we are in agreement that the first two pillars can't make any money, but what about the final structural Blockchain pillar? It leaves NFTs as the only monetizable component of Wyplay Community, through revenue sharing agreements with operators selling exclusive NFTs. As expected, Feral admits that Wyplay's share will be minimal.

We press for an estimate. Less than 10%? Less than 5%? He smirks awkwardly, telling us he isn't going to reveal numbers, as Wyplay's slice of the pie remains a detail yet to be ironed out with operators. Of course, Community doesn't have any live customers yet, although with all the positive feedback Wyplay has supposedly received from the marketplace, some must have expressed a ballpark estimate of what they would be willing to pay Wyplay for each NFT transaction.

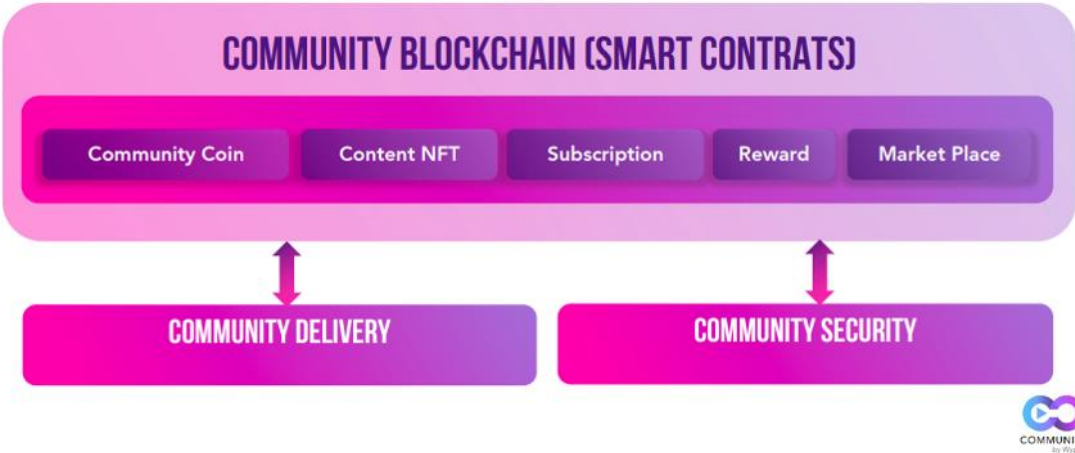
Whatever it might be, it's small, really small, so Wyplay is relying on huge scale NFT transactions for Community to become profitable.

Similarly, Feral isn't budging with how much Wyplay has invested upfront in Community. It took a dedicated team of 15 people a full year to develop Wyplay Community, and it still isn't finished. Our

**No one makes money from DRM, or P2P – meaning all hope lies in nascent NFTs**

estimate is at least \$1 million to date, potentially doubling as the platform adds more business rules to its smart contracts from one customer to another, adapting its blockchain connectors to existing back-end systems (see image).

For us, this is the most significant technical aspect of Wyplay Community. It has to plug-in to existing back-end systems with zero disruption to an operator’s video technology stack or network infrastructure.



Wyplay’s roots in middleware have taught it a lot about David versus Goliath situations. Its flagship Frog middleware jumped into the market around 2014 with early success, but was battered by the

arrival of Android TV Operator Tier and RDK that took pay TV platforms by storm, through standardizing functions in the video software stack, including middleware.

Frog couldn’t compete, but it could use these new technologies to its advantage, implementing elements of Frog within these standardized stacks for operators – becoming agnostic while guiding customers to take the best of both worlds.

Wyplay has done that pretty successfully since, with Frog accounting for around 70% of total revenues today.

Returning to the blockchain discussion, our patience begins to wear thin. Wyplay Community is primarily targeted at large pay TV operators with millions of customers (Feral pulls an average of 4 million from nowhere), most of which are of a certain age (Canal+ has an average viewership of 53 years old, for example).

Wyplay’s thinking is that pay TV operators need to differentiate and attract younger viewers in order to continue operating profitable TV businesses, and that NFTs are the way to do that. We argue that most operators can’t even successfully transition from cable/satellite TV delivery to IP-delivered video to attract younger generations away from cheaper, or free, alternatives – let alone springboard into monetizing cryptocurrencies.

It hinges entirely on an operator's individual rewards scheme. If the rewards are compelling enough, then you will have a credible business model. If not, you will be laughed out of the blockchain. Feral provides a handful of examples.

If a subscriber agrees to reshare NFTs to increase royalties, they will receive a coin-based reward. If they promote it on socials, they get coin. If they help to bring in new subscribers = coin.

Ultimately, Wyplay Community has to become more customizable around an operator's needs. From where we're sitting, the issue is that operators don't know the first thing about what young people want – throwing money at walls and hoping it sticks for more than six months at a time.

It is all well and good preaching about reducing carbon footprints using decentralized infrastructure, but the business models we are hearing from the industry are murky at best.

Wyplay is so focused on NFTs, where merchandise and live events are the primary pull, that it hasn't even considered how to monetize ads in the blockchain, which for many video clients looks to be a more promising opportunity.

It's a big gamble from Wyplay, despite Feral's protestations that his company has made the right choices, followed up with the ominous caveat of, "We'll see how it progresses."

## **Triple and Quad Play**

### **RDK-B bolts on Airties' EasyMesh, as telcos consider common stack**

Ask and you shall receive.

Six weeks ago, Faultline piled pressure on RDK Management – proposing that the group was stalling on supporting EasyMesh as a feature of RDK-B, a year after the WiFi Alliance's mesh standard was added to the RDK-B roadmap.

We observed how tier 1 operators and WiFi technology vendors alike have been showing signs of being ready to embrace EasyMesh within next-generation router software stacks, yet there were no signs from the RDK camp of EasyMesh integration being in motion, until now.

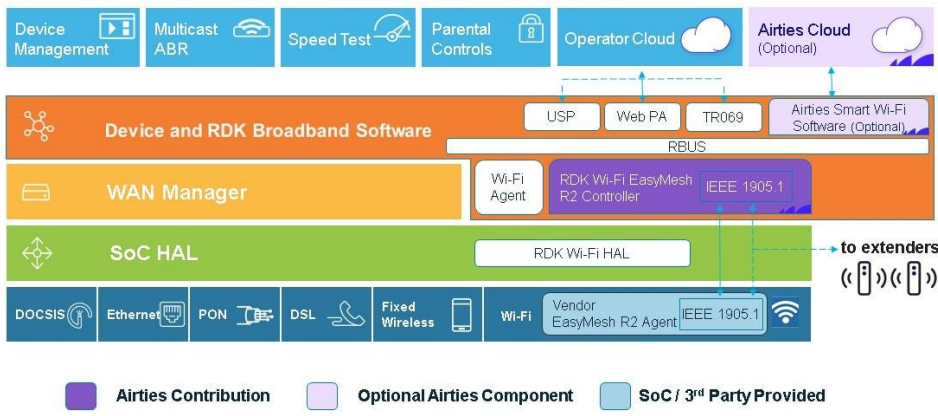
The radio silence has been ended by Turkish WiFi management specialist Airties, which has contributed its WiFi EasyMesh controller software module as a new component for use by the RDK community across broadband gateways.

For internet service providers, it means less time spent plugging away on laborious platform integration efforts across OEMs and chipsets – time that internal teams can instead channel into more proactive activities. In Airties’ case, it wants ISPs to spend more time developing value-added managed WiFi services for subscribers, for obvious reasons.

The first is that most operators will eventually standardize on EasyMesh APIs anyway, so you might as well embrace it, as a translator for WiFi components and access points (APs) to create a common framework within a mesh WiFi architecture.

In turn, this means shrinking demand for proprietary mesh architectures, like Airties’ own, and more room for controller-initiated steering, cloud-based channel management, and cloud-based diagnostics collection.

As the attached graphic shows, money is to be made from the optional Airties Smart WiFi software located in the layer above its controller contribution, and a further layer above still for the optional Airties Cloud.



Here, channel and band steering is handled using what Airties describes as a unique hybrid cloud-edge architecture – based on real-time network conditions, embedded intelligence in CPE, and Airties’ cloud management capabilities.

The controller can support any certified WiFi EasyMesh device.

Within the core functionality of the new RDK-B component, contributed by Airties, you will find support for onboarding WiFi EasyMesh agents on gateways and extenders; setting up in-home backhaul connections; propagating fronthaul configuration changes; and collecting statistics periodically. Furthermore, a pending enhancement will expose configurations and device state via the common RDK message bus (RBus).



For vendors, they are now free to deploy their own products using this contributed WiFi EasyMesh controller – adopting RDK-B’s WiFi data elements and multi-AP data model for the EasyMesh controller.

Looking long-term, what is arguably more significant here, as we gaze into our crystal ball, is the issue of fixed-mobile network convergence. Tier 1 European operators Deutsche Telekom, Vodafone, and BT have recently declared their support for an open source objective to provide a production grade 5G stack capable of integration with RDK-B and OpenWrt frameworks. OpenWrt is an open source project for embedded operating systems on wireless routers, based on Linux.

With growing interest among operators in developing a common stack to facilitate consolidation of core and aggregation networks, this has led to the development of an open source standard called WWC Reference Implementation for 5G-RG (residential gateways), designed to allow fixed and mobile services to be consolidated over a common core.

With Airties in the process of being acquired by private equity firm Providence Partners, we deliberated last week how the company’s future could hinge on appealing more to the mobile guys. Faultline is currently in the process of setting up a call with Airties leadership to extract extra details on the acquisition, RDK-B, emerging markets, and its possible mobile-influenced wardrobe change.

## **Triple and Quad Play**

### **Dish Network defying financial odds, despite churning TV subs for fun**

Only AT&T and Comcast lost more pay TV subscribers in Q1 2022 than Dish Network’s 462,000. Of course, Dish’s damage is split across two video products – the Dish TV satellite offering and the streaming service Sling TV – making the losses of 228,000 on Dish TV and 234,000 on Sling TV appear trivial when split, compared to Comcast’s 484,000 and AT&T’s 620,000.

While last week’s pay TV analysis from Faultline highlighted how Comcast’s cable video revenue for the quarter was relatively stable, at a loss of just 1.5% to \$5.54 billion, Dish Network’s pay TV revenue looks even more remarkable, all things considering. The satellite TV operator saw Q1 2022 revenue fall just 1.1% to \$3.13 billion, with a positive operating income of \$752 million, despite having almost half a million fewer customers.

Looking at the wireless business for comparison, where Dish Network is channeling increasing focus and capex, revenue sunk 9.5% year on year to \$954.4 million, at an operating loss of \$201.7 million.

Worse, it lost another 343,000 mobile subscribers, bringing the total down to 8.2 million, on the eve of its 5G network expansion, with Dish having just five weeks left to meet the FCC's construction deadlines. Headlines paint a very different picture.

Our point is that, somehow, select US pay TV heavyweights are weathering the cord cutting storm, as sustained subscriber losses fail to reflect on finances. ARPU is helping here in a significant way, with Dish Network's average pay TV subscriber paying almost \$11 more today compared to two years ago – rising to \$99.44 as of Q1 2022 across both Dish TV and Sling TV.

ARPU breaching the \$100 a month barrier across both Dish TV and Sling TV is almost unthinkable, given the wealth of cheaper streaming alternatives on the market today.

**Somehow, select US pay TV heavyweights are weathering the storm**

But with the average US household stacking four OTT video services, and many signing up to five or more video streaming offerings, many of Dish's remaining 10.25 million pay TV subscribers (7.99 million Dish TV; 2.25 million Sling TV) are happy to pay a little extra for everything all in one place, with live TV as the focal point.

Dish's wins in ARPU and resultantly revenues could therefore be interpreted as a content aggregation success story, and a testament to some of the newer advancements to its flagship fleet of Hopper DVRs, such as betting features and the Android TV Hopper Plus platform, which landed recently.

The Android TV-based platform features the ad-skipping functionality that is a trademark of the Hopper set tops, and Dish is marketing the Android TV version as a more of a "whole home" experience compared to a home that might have a Roku in one room and a Chromecast or a Fios TV from Verizon in another room.

Hopping ads continues to be a strong strategy for Dish's TV business because, apparently, the biggest complaint the operator receives for linear TV is having to sit through some 15 minutes of commercials per hour of TV. We find it hard to believe that Dish Network's call centers are stacked with queues of disgruntled ad viewers, over those complaining about subscription fee increases, billing problems, or technical issues.

Speaking during the company's earnings call, Dish Network's decisionmakers discussed "feeling around" the market to figure out how it can continue to improve the user experience of its video offerings while accepting that advertising makes a lot of money.

Co-founder and Chairman Charlie Ergen believes the video business will be "very profitable" going forwards. Playing devil's advocate, we would argue that Dish Network has no right to be celebrating any successes given its monstrous video subscriber churn rate, although – as we said – the finances look surprisingly strong.

Churn has accelerated since the height of the pandemic, when pay TV operators enjoyed relatively low rates of subscriber turnover compared to trends of previous years. That period of respite is over, one that Dish management believe provided valuable time to concentrate on retaining the right type of customer and – most importantly – ensuring that these select customers are profitable.

The operator acknowledges that TV bills are too high, but at the same time consumers are willing to pay for the thrill, reliability, and convenience of live TV. This ties into Dish's strategy of targeting more rural customers, which it has been pushing since around 2016, so nothing new here, but there are signs of this paying off.

Dish is still not willing to disclose its rural penetration, but does at least acknowledge publicly that the majority of its TV customers are situated in rural areas of the US.

Meanwhile, the conversation around Dish Network's merger with DirecTV is not going away. Ergen continues to describe a coming together of the two businesses as "inevitable".

## **Advertising**

### **Full Throttle scales on coattails of desperate MVPDs**

With everyone shouting about the death of the cookie, those with a genuine solution to the problem can find it hard to be heard above the white noise. First-party data miner Full Throttle has an ingenious possible answer with potential for huge scalability in the coming years – by bundling the technology into the Upfronts pitch of leading MVPDs.

Publishers like Charter, Comcast and iHeartRadio are flashing a 12-month free subscription to Full Throttle's SaaS service in return for an annual ad commitment at the Upfronts.

Speaking to Faultline this week, the company's CPO, Amol Waishampayan, says that despite the fuss over a cookie-less future, advertisers and brands are facing a huge challenge of scarcity when it comes to first-party data. At present, brands are scraping the barrel for email addresses of users that never followed through with a purchase, but Waishampayan claims this is only the tip of the iceberg.

Full Throttle looks to generate unique, cookie-less first-party data for its brand partners, which can then be used across all advertising verticals. This is done without using any of the standardized identifier technologies (e.g. cookies, IDFA), most of which are heading for the chopping block in the near future.

Instead, Full Throttle installs its proprietary, patent-pending software directly onto a brand's website. Subject to them opting in, this tracks a user's activity on the website and their location data, which means that this previously untraceable consumer can now be linked to an address for retargeting. Waishampayan says that this offers a ten-fold increase in the amount of household-level, first-party data that a brand has access to.

**"It's not our data – we are just the service provider"**

"It's not our data. We are just the service provider, allowing marketers to access their own stores of data," Waishampayan explained.

For the moment, this allows brands to perform "retargeting on steroids" – essentially, chasing up all those potential customers with follow up ads. But as well-loved identifiers begin to die, these proprietary data stores will become essential for all marketing activities.

On the video side of things, Full Throttle can distribute a brand's collected data to MVPDs via the decentralized media network Blockgraph. In the US, Blockgraph currently spans over two-thirds of all broadband and pay TV homes. Our chat with Blockgraph's CEO, Jason Manningham, just a few weeks ago, teased that the company is set to add another 19 million pay TV homes in the coming months as it adds more MVPD partners.

This should be a significant leg-up for Full Throttle's expansion, although Waishampayan cautioned that fully onboarding onto Blockgraph is a lengthy process for MVPDs. Operator engineers have to facilitate an entire new data flow to and from Blockgraph, which means that no transactions will be taking place right away.

Brands can also plug their sales data into Full Throttle's dashboard via APIs to use the new sets of first-party data to establish whether campaigns led to any conversions.



Waishampayan says that many advertisers are growing weary of huge data marketplaces like LiveRamp, which are increasingly seen as a black box. “Everyone used these because of great brand equity but increasing demands of transparency from TV networks means that advertisers often don’t trust them,” he explained.

Another issue is the fact that LiveRamp does not offer true ownership of the data that fuels its identity graph, which Waishampayan believes could prove problematic in the long run. “We are evangelical about enabling people to own their datasets,” he said.

This trepidation about future-proofing its business means that Full Throttle is thinking outside the box. Waishampayan says that enabling advertisers to safely cross-reference their datasets, without facilitating actual sale of the data, is next on the agenda, although he cautions that this is still at least one year out from coming to market.

This would mean that every brand or advertiser could become a publisher, monetizing datasets not through sale, but through exchange. “The best way to be resilient is not just selling anonymized data – it will still count as a sale,” Waishampayan warns.

## **Advertising**

### **Strategus guides brands through CTV ad maze, trashes walled gardens**

As we cover video ad tech increasingly, Faultline cannot help but get a sense of déjà vu. The endless seesaw between fragmentation scares and empty standardization promises can often leave us feeling seasick, but this also creates opportunities for vendors.

One such company holding advertisers’ hands is Strategus, a managed service provider that helps brands and agencies to run and measure highly targeted connected TV (CTV) campaigns.

While digital advertisers are largely well served when it comes to single self-serve ad platforms – think Google or Facebook’s in-house ad tools – there is no such comparable service for CTV. Speaking to Faultline this week, the company’s Co-founder and SVP of Strategy and Innovation, Joel Cox, says this fragmentation makes a “white glove managed service” like Strategus essential.

“Google and Facebook have been uniquely successful at simplifying a highly complex advertising process, but that same privilege does not apply to CTV,” Cox explained.

Since launching the first programmatic CTV campaign back in 2015, Strategus has looked to lead the way as traditional TV advertising budgets trickle over to CTV. “Marketers just don’t have the time to learn how to use cutting edge ad tech for CTV,” Cox says.

Strategus works with a range of brands across retail, franchise, and direct-to-consumer verticals – everything from car dealerships to breweries. These clients come with a campaign goal, ad creative and a target audience, and it is Strategus’ job to find the CTV audiences and deliver the results.

“The traditional mantra for video advertising has always been ‘position and placement,’ but that just doesn’t get us a fraction of the way around the track these days, although many of the linear guys still stick to it,” adds Strategus President Todd Porch, also on the call, who previously held significant advertising stints at both Comcast and Yahoo.

Cox concurred that effective targeting relies upon knowing exactly who is watching a title, rather than working from the title they are watching.

To achieve this, Strategus weaves together a range of siloed data sources to glean exactly who a CTV user is for advertisers. Location data can be taken by the likes of Foursquare and Factual, while online shopping data can come from Oracle.

With Strategus reliant on open data ecosystems, it was hardly a surprise to hear that Cox does not approve of certain CTV vendors going for the walled garden approach. He pointed to Vizio’s Inscape ad tech arm, which sells user automatic content recognition (ACR) data to exclusive partners – a move that saw the US TV maker’s share price drop from \$27 to a current trading price of around \$7.

Cox believes the market knows open data ecosystems are the way forward for guaranteeing premium ad experiences, and that eventually the pendulum will swing back. “People will try to sell their own data, but the ad pods will not sell, and users will churn. Eventually those companies will come back to an open inventory,” he predicted.

Cox pointed to a similar pattern that emerged in display advertising in the early 2010s. “Every website tried to sell their own inventory, but they quickly realized they did not have enough tentacles to reach the right advertisers and audiences,” he continued.

**“Position and placement doesn’t get us a fraction of the way around the track these days”**

Strategus also does much of the legwork when it comes to campaign analytics and performance measurement, joining the dots to track all consumer actions post-ad exposure. Like everything in ad tech, Cox again expressed that these data sources are not only fragmenting but getting deeper and wider.

“People assume this industry is ‘plug and play,’ but there is just so much complexity,” Cox explained. “We are bringing the simplicity.”

So fragmented is the road to a successful CTV campaign that Cox has his fair share of horror stories. One well know brand accidentally spent \$1 million on CTV ads over just one weekend, simply by entering a few values incorrectly to a demand-side platform (DSP). “They aren’t getting their money back,” he warned.

Cox points to Netflix’s recent stall in subscriber growth as a sign that SVoD is in decline. While a landmark moment in OTT video, Faultline is hesitant to jump to such grand conclusions. The 0.09% loss in subscribers can be neatly drawn to sanctions on Russia, while our sister research wing Rethink TV forecasts that Netflix is still set for small growth in North America through to 2027.

Nonetheless, Cox says that as ad-supported streaming grows, platforms need to start treating ad tech and experiences with the same reverence that they treat content. “Platforms should have fewer ad pods, they should be highly targeted, and appear in low frequency,” Cox demanded.

Cox believes that fewer, better targeted ad pods are the way forward, as they deliver more yield from an advertisers spend compared to a traditional ‘spray and pray’ campaign. Strategus can get granular across its national US footprint, with targeting capabilities going from large DMAs down to certain ZIP codes.

Faultline could not resist alerting Cox and Porch to the persistent AVoD bee in our bonnet – the fact that most pure-play AVoD platforms still lack the premium content that premium advertisers desire. Cox says marketers should remove all assumptions about certain content being seen as ‘lesser’ when advertising. If targeting is done right, you will still reach who you need to get to.

In terms of inventory, with the forgivable exception of YouTube, the company serves all the big names in CTV advertising of Roku, Tubi, Peacock, Crackle Plus, Vevo, Twitch and Crunchyroll. “Curating great inventory is easy. We excel in data, analytics and attribution,” says Porch.

**“People assume this industry is ‘plug and play,’ but there is just so much complexity”**

## Sustainability

### Tachyum tackles idle waste with reprogrammable Prodigy silicon

FPGA specialist Tachyum has made an outlandish claim this week, by declaring its new Prodigy line of reprogrammable silicon as the “first universal processor.” With big performance claims to boot, the Prodigy FPGA is being positioned to target all manner of data center workloads, with video being a key use case.

Of course, Faultline is pretty familiar with such grand claims, and we are used to sifting through some very dubious press releases. Tachyum, thankfully, does not fall into that camp. CEO and co-founder Radoslav Danilak talked us through the road that led to the Prodigy.

The concept behind the universal processor is one that can properly handle the three main silicon approaches needed in enterprise workloads – CPU, GPU, and AI-focused designs like Google’s TPU (Tensor Processing Unit).

Danilak stressed that the Prodigy is not simply three technologies in a trench coat, and the key feature is the ability to repurpose one of these chips to match the dynamic workloads seen throughout the day.

On this point, Danilak highlighted the problem that these idle servers pose. Data centers, according to Tachyum’s research, represent 4% of total global power consumption, but this is a figure that is growing at 15% per year – and faster during Covid. At this pace, it would reach 40% by 2040, and so something needs to be done.

Worse, much of this power is being spent on servers doing nothing. Facebook recently published data showing that the average utilization of its servers over a 24-hour period is around 30% - mimicking human activity cycles. Due to latency issues, you cannot simply run a global data center, and serve EU customers with a US rack.

The obvious thing to do here is to use the idling servers for something that is productive.

In the Facebook example, Danilak pointed to how these servers should be training the AI models that Facebook increasingly relies on, but here is where the differences in silicon architecture come in – these servers are not physically suited for that sort of job.

And so, the pitch is that a product like Prodigy would allow companies to actually get the most out of their investments – either for internal jobs, or providing third-parties with compute power. Danilak pointed to Microsoft selling surplus Azure AI cycles to the US government, for video surveillance processing.

The scale of the problem seems enormous. Danilak said there is around \$300 billion of commissioned IT equipment currently, and at that 30% utilization rate, this effectively means \$180 billion has been wasted. Here, Tachyum proclaims that using its chips would allow you to avoid that headache, and save on both cost and power consumption – as idle servers are still consuming rather a lot of electricity.

The key difference in Prodigy to the CPU, GPU, and TPU rivals is that Tachyum spent a lot of time developing a different way to move data around inside the chip – a new instruction set, effectively.

Over the past few decades, there have been immense improvements in semiconductor performance, with the raw speed of transistors improving between 6x and 8x. Moving more elements from the PCB inside the chips themselves has also cut latency significantly, but the industry has been grappling with a slowdown in Moore’s Law simply due to the rules of physics.

**The scale of the problem seems enormous**

As the process size of these chips has decreased, the internal wire size has too. This unfortunately means that the electrical resistance of the wire has increased, which put simply means that a 10x smaller wire ends up being 100x slower.

Danilak said that these wire delays are now limiting the performance of functional blocks,” and that “as time has moved on, more processing time is spent in the wire than the transistor.”

This forced the company to create a different approach for the internal operations.

“The big question was if we could break the process of moving the data around and the calculations into two different processes,” said Danilak.

“This allows us to do the next calculations at the same unit that just produced the data, without having to move data around the wires. That’s the simple way to explain it. We are moving data around the 8 positions in the unit in a way that is faster in about 93% of cases. We can’t solve the physics problem, so have to find a way around it,” Danilak continues.

Danilak added that if you are not moving data around, you are also not using electricity to do so. This leads to a claimed 3x to 10x improvement in power consumption. In the variety of performance throughput benchmarks shown to Faultline, the lowest ratio was 3.3x, with the highest being 17.1x.

Danilak stressed that 6x to 8x is a sensible expectation, with a rack-level comparison showing one air-cooled server rack of Prodigy chips scoring 4.8x the performance of rivals. Liquid cooling improves this further, often doubling it, and the best scenario saw one rack of these Prodigies do the same work as 13 Nvidia DGX racks. Liquid cooling also allows for energy recovery projects like district heating, which would allow data centers to offset their carbon emissions further.

Pricing is not finalized, and this was the only time that Danilak did not have a concrete answer for us. Instead of providing a price, he discussed the cost of these chips compared to rivals. Tachyum is using TSMC and Samsung to produce these 5nm chips, and Danilak said that between \$3k and \$4k of Prodigies should provide the performance of \$10k in rivals.

**This was the only time Danilak did not have a concrete answer for us**

Pricing will be confirmed later in the year, once production volumes give a clearer cost expectation. Sampling is scheduled for the second half of the year, and the second generation of Prodigy chips is expected in 2024. Tachyum is working on a total addressable market (TAM) estimate of \$3 billion to \$4 billion, saying that the AI-based market is around 10% of the \$30 billion spent on microprocessors for the cloud. The percentage will increase in time.

Currently, Tachyum is helping to assemble servers, and has its own motherboard reference design to aid in this. Eventually, Tachyum wants to sell just the chips, and once its server partners no longer need Tachyum reference designs, Tachyum will stop making its own servers.

The Prodigy family spans from a 32-core chip to a 128-core model, designed to cater for all data center budgets. The 128-core variant has a monstrous 5.7 GHz clock speed, supports 32 TB of RAM, with a 200 Gbps Broadcom Ethernet interface, and the architecture lets customers run native x86, ARM, and RISC-V.

Tachyum has invested heavily in creating the software ecosystem to support its new instruction set too, noted Danilak. In terms of physical size, the chip is around 500mm<sup>2</sup>, which is significantly smaller than an Nvidia H100, at about 800mm<sup>2</sup>, and Intel's leading Xeons, which measure some 650mm<sup>2</sup>.

The team which includes Steve Furber and Fred Weber, who Danilak described as the father of the ARM instruction set and the person that developed AMD's 64-bit x86 instruction set respectively, the latter of which was so much better than Intel's attempt that it continues to license it to this day. "Having two people responsible for 98% of computing running today in Tachyum is pretty amazing," said the CEO.

## **Sustainability**

### **Meta ruffles OpenAI feathers with carbon-crunching language models**

A research paper hot off the press from Meta AI has made quite a splash in the academic community. In layman's terms, Meta has effectively open sourced a set of language models so large that it makes Meta more open than OpenAI, the dedicated research lab founded by Elon Musk, among others.

While it sounds like Meta is ticking all the right boxes – including sustainability, diversity and openness within AI language models – one has to consider whether this is more than a muscle-flexing PR stunt? Besides, what's stopping people using Meta's new open models for harmful purposes, like spreading fake news or developing more convincing deep fake videos?

Without getting too deep into the methodology (you'll thank us later), Meta's open pre-trained transformer (OPT) large language models (LLMs) bear potential to make strides in fields including modern dialog models, bias and toxicity evaluations, and hate speech detection.

With the revolution of social media as a significant political weapon, and with grave concerns about how young people will interact with future metaverse experiences, making these LLMs available to the wider research community is designed to bring benefits to AI and natural language processing development. However, we doubt the Facebook research wing will receive the plaudits it deserves for this huge open source effort.

But why bother trying to get one over on OpenAI in the first place? The reason Meta's OPT has caused such a stir is that the suite of decoder-only pre-trained transformers range from 125 million to 175 billion parameters, which is comparable to the 175 billion parameters comprising OpenAI's GPT-3, but at 1/7<sup>th</sup> the carbon footprint to develop.

Called OPT-175B, Meta achieved this using the latest generation of Nvidia hardware.

Models of this size require massive computational resources and therefore cost, meaning any copycat attempts need serious capital and come at sustained detriment to the environment. Meta warns that repeated efforts to replicate a model of this size will only amplify the growing compute footprint of these LLMs, which sounds a little hypocritical, although at 1/7<sup>th</sup> the carbon footprint, it's hard to argue.

In other words, "Don't even bother trying to beat us because you'll fail, and we'll guilt trip you for impacting the environment even if you do produce better results than us." Meta has released its logbook detailing the infrastructure challenges faced, as well as code for experimenting with all released models.

Given the centrality of LLMs in many downstream language applications, Meta also hopes to increase the diversity of voices defining the ethical considerations of such technologies. These voices across the AI community span academic researchers, civil society, policymakers, and private sector industries.

For the full OPT-175B suite, research access is provided upon request, although the models between 125M and 30B parameters will be released without request.

As for limitations, there are several, including of particular concern that OPT-175B can produce factually incorrect statements which can be harmful in applications where information accuracy is critical. OPT-175B was also found to have a high propensity to generate toxic language and reinforce harmful stereotypes, as well as a tendency to be repetitive and easily get stuck in a loop. It is within the limitations that it makes sense to open source the LLMs and where arguably most research efforts are required.

## Worth Noting

### OTT Video News, Deals, Launches and Products

#### Five years ago this week...

**Disney** was suffering from buyer's remorse, shunting its multi-channel network (MCN) **Maker Studios** off to its in-house branded content division, Disney Co/Op to assist with the made-up job of "influencer marketing". Following an impulsive spending spree in the early 2010s, legacy media firms were coming to terms with the fact that MCNs and their content creators were not as lucrative as first hoped. **Disney** had paid over \$500 million for Maker Studios, but its interest in the MCN quickly waned, with the number of the creators on the roster slimmed down from 60,000 at the time of purchase, to just 300.



--

**Warner Bros. Discovery** will form a 50:50 joint venture with **British Telecom**, whereby the **BT Sport** business will be transferred to WBD and merged with **Eurosport UK**, which WBD owns. Sports streaming service **DAZN** had its bid to acquire BT Sport rejected by the UK operator.

**Faultline** has been openly critical of companies like **NBCUniversal** when it comes to empty promises of revolutionizing how video advertising is consumed through interactive experiences. The **Comcast**-owned company finally has something to show for all its hot air, including the smartphone-based AR Portal as a pathway to advertising in the metaverse, without (by the sound of it) requiring a headset. NBCU is also promoting sequential storytelling as a highly targeted ad play, using first-party data to target consumers on three separate occasions throughout the day.

**Liberty Global** closed Q1 2022 with just under 4.5 million video subscribers spread across the five territories of Belgium, Switzerland, Ireland, Slovakia, and the Netherlands. All but Switzerland saw subs losses in the quarter, with the **VodafoneZiggo** Dutch joint venture hardest hit at minus 70,800. Since the merger of the **Virgin Media** UK operation with **O2**, video subscribers are no longer broken out, while the Polish **UPC** operation has been discontinued following the sale to **Iliad's Play** last month, when it had 588,200 video subscribers.

US DVD rental company **Redbox** has been bought by **Chicken Soup for the Soul Entertainment**, for \$375 million. Redbox completed its IPO in October 2021 at a valuation of nearly \$700 million, marking a dramatic climbdown.

Public broadcaster **Channel 4** has taken a step towards privatization through a landmark deal with **YouTube**, amid ongoing outrage about a pending sale by the UK government. Hundreds of hours' of Channel 4 content will be injected into the video sharing platform in such a way that allows Channel 4 to sell its own advertising, which is a first-of-its-kind move in Europe.

Another SVoD platform is moving in on AVoD and FAST, with **Alchimie** seeking to strengthen its European footprint through an ads partnership with **LG Channels**. The deal makes the Alchimie Channels OTT video offering available on LG Smart TVs running webOS 3.0.

**Xperi**, the parent company of **TiVo**, reported Q1 2022 revenues up 16% to \$257.4 million. Media business highlights from the quarter include a long-term patent portfolio license with a top 10 vMVPD, a new win with US cable company **NfinityLink Communications**, and double-digit subscriber growth in IPTV. On the media product side, Xperi credits the integration of **YouTube TV** into TiVo Stream OS and 4K, plus the launch of the TiVo Xtend for connected TV advertisers to grow reach, as well as advancing the TiVo Stream ecosystem with new services, content partners, OEMs and chipset makers.

**M7 Group** is jumping headfirst into smart TVs via the HbbTV OpApp specification in Europe, complying with the M7 Conditional Access Module and integrating Fast Scan Tables. From June 2022, M7 channels will be available on any smart TVs running **Mediatek** chipsets, starting with **TP Vision**, without any additional hardware required.

**Vivendi**-owned French pay TV operator **Canal+** is reportedly looking to acquire US TV network **Starz**, following rumors that **Roku** is in the running, as **Lionsgate** looks to divest the business.

**AMC Networks** has 9.5 million total subscribers across its video streaming properties after gaining 500,000 users during Q1 2022, although it doesn't say how many are on the flagship AMC+ platform. New guidance puts projected total streaming subs at between 20 million and 25 million by 2025.

**Telefónica**-owned Spanish operator **TSA** has deployed monitoring tools in its head-end from Swedish firm **Agama**. A full transport stream analysis is covered by multicast analyzers, widgets, and dashboards, with additional network information provided in a single customizable view.

**Akamai** has released what it is calling the world's first NFT (non-fungible token) artwork that is dynamically powered by the internet. This piece of marketing fluff from the CDN heavyweight is described as a one-of-a-kind digital artwork projecting the world's internet activity, via Akamai's 4,000 PoPs, that visually changes in real-time based on internet traffic changes and cyberattacks. Final bids are due June 24, 2022.

**Roku** is polishing its NewFronts spiel through a new partnership with **Microsoft**, so that advertisers can draw links between ad exposure and subsequent internet searches, using Microsoft's audience intelligence. Roku is also offering dynamic linear ad insertion, expanded measurement, clean room data, and a Roku advertising watermark.

Switzerland is holding a referendum on May 15 to decide whether streaming services should be forced to invest in local content production. Under the proposals, top tier SVoDs would have to invest 4% of revenues into local content.

India lost 370,000 DTH subscribers in Q4 2021, according to figures published by the **Telecom Regulatory Authority of India**, with the total number now at 68.52 million homes. **Tata Sky** has the largest market share of DTH at 33.48%, followed by **Bharti Telemedia** (26.37%), **Dish TV India** (22.04%) and **Sun Direct TV** (18.11%).

The merger of Canada's **Rogers** and **Shaw Communications** might not go ahead, despite approval from both regulatory bodies. The Commissioner of Competition is now filing an application to oppose the merger, due to worries of its impact on Canada's competitive wireless market.

**Facebook Gaming** and **YouTube Gaming** have suffered viewership drops in Q4, according to new figures from **Streamlabs** and **Stream Hatchet**. **Twitch** managed to grow its in Q4, but the three services combined for an 8% drop in viewing hours year-on-year, from 8.8 billion hours to just over 8 billion. YouTube's decline is impressive, given its attempt to poach talent from Twitch over the past year.

**Liquid Intelligent Technologies**, the pan-African technology group that is attempting to accelerate broadband adoption on the continent, has acquired **Telrad Networks**, an Israeli telecommunication equipment manufacturer. Telrad seems to specialize in fixed wireless offerings, and was bought for \$58 million (199 million NIS) – a big slump from its IPO valuation of \$99 million.

**Frontier Communications** has paid the **FTC** \$8.5 million to settle an investigation into misleading consumers with broadband speed claims. Frontier is not admitting guilt, and the fine is essentially equivalent to a full day's profit. This is not the best way to emerge from Chapter 11 bankruptcy, either.

**Microsoft** has brought **Epic Games' Fortnite** back to iOS, via the Xbox Cloud Gaming service. Microsoft would not waste an opportunity to turn the screw on rival **Apple**, but for Epic, the availability looks set to mute its continued attempts to challenge Apple's market position in the courts. A single jury decision on the role of the mobile platforms as gatekeepers to games and video services could cause immense upheaval in the sector. Meanwhile, Microsoft is expected to launch smart TV apps for the cloud gaming service this year.

**TelevisaUnivision** has acquired **Pantaya**, a Spanish-language streaming service with some 900,000 subscribers, ahead of the global launch of TelevisaUnivision's ViX+ service later this year. No price has been given, but Pantaya was originally a joint venture between **Lionsgate** and **Hemisphere**. Lionsgate sold its majority stake last year for \$124 million.

**intoPIX**, **Nextera Video**, and **Adeas** have partnered up to show off a reference design for the new IPMX AVoIP (AV over IP) standard. intoPIX's hardware-focused video compression and video-over-ethernet kit will be used by Nextera and Adeas in FPGA implementations, for moving a 4Kp60 4:4:4 video signal over Cat 5e cables. That the demo will be shown at the **AMD Xilinx** booth clues us in to the silicon provider.

**Faultline** - Analyzing disruption in media, broadband and advertising - for over 20 years.

Faultline is a merger of 3 industry-renowned publications: Rethink's Faultline, with Rider Research's Online Reporter and Internet TV Reporter. The combined result is a weekly service spanning topics including OTT video technologies, broadband, WiFi, devices, advertising, content, and video formats - studying disruptive emerging market trends.

Faultline is published by Rethink Technology Research Ltd, S4, Temple 1852, Lower Station Approach, Temple Meads, Bristol, BS1 6QS Tel: +44 (0) 117 925 7019  
Copyright Rethink Technology Research Ltd 2022

**Faultline** Editor: Tommy Flanagan [tommy@rethinkresearch.biz](mailto:tommy@rethinkresearch.biz)  
Senior Contributor: Alex Davies [alex@rethinkresearch.biz](mailto:alex@rethinkresearch.biz)  
Contributor: Rafi Cohen [rafi@rethinkresearch.biz](mailto:rafi@rethinkresearch.biz)

### **About Rethink**

Rethink is a thought leader in quadruple play and emerging wireless technologies. It offers consulting, advisory services, research papers, plus three weekly research services; **Wireless Watch** which has become a major influence among leading wireless operators and equipment makers; **Faultline**, which tracks disruption in the video ecosystem and has become required reading for anyone operating in and around quad and triple play services and digital media.

### **Distribution contact details:**

Roz Hilton—[roz@rethinkresearch.biz](mailto:roz@rethinkresearch.biz)  
+44 1 962 732 886  
+44 7 399 491 840

### **Editorial contact details:**

Tommy Flanagan—[tommy@rethinkresearch.biz](mailto:tommy@rethinkresearch.biz)  
+44 (0) 117 925 7019