Operator: Thank you everyone for joining us today. This is the first session in what we hope will become an ongoing educational webinar series for investors, where we periodically discuss industry-relevant themes and developments. Today's discussion is strictly focused on AI in media, so we request all participants to keep their questions limited to this topic. This session will run for 1 hour. Bhaskar will speak for the first 40-45 minutes, followed by Q&A in the last 10 minutes.

As this is a listen-only webinar, please type your questions in the chat box with your name and company. I will voice them during the Q&A. At the end, you will be invited to share feedback through a Google form by scanning a QR code. Please take a moment to share your inputs and help us improve future sessions.

Before we begin, please note the disclaimer. Questions relating to the company's financial performance, business outlook, operational metrics, strategy, or other company-specific matters will not be addressed. Please note that certain comments may include forward-looking statements, industry observations, or estimates based on management's current views and assumptions. Actual outcomes may differ materially. With that, I'll turn it over to Bhaskar to kick us off.

Management: Thank you very much, Amulya. Thank you everyone for joining us. Good morning, good evening, and good afternoon, wherever you are. Today's topic is an interesting and exciting one: AI in media. I am sure AI is a topic we are all working on across many different industries, and media is no different. We are going to talk about it today.

If we look at it, the first thing I wanted to talk about is to understand how a typical video value chain works. As you know, for television and all the OTT video streaming that we see, it starts from the camera and ends with the screen where we watch content. We watch content on television sets, phones, or tablets. Wherever you are watching, that is where the content comes through, but it always starts with the camera. There are several sections to how the process works today.

For example, let's look at a game. If you are watching a cricket match, the action obviously starts in the stadium. At the stadium, the first thing you see is that you have multiple cameras. All these cameras get connected to an OB van standing outside the stadium, typically with a satellite antenna on top, through which it gets beamed to a studio. Anyone who has gone to a stadium would have seen this mechanism. There are commentary boxes from where commentary is being done, and there is a production room where they decide which camera content makes sense and perform the switching of the cameras.

Once that is done, it comes to a preparation stage, typically a studio. In the studio, this cricket match is converted with sports scorecards, logos of the different teams, multiple audio tracks, and of course, the advertising in ad breaks. All of that gets added in the preparation stage in a typical studio environment.

Once that is complete, it gets distributed, either through cable or through streaming TV platforms, through which we watch this content. This is typically a stage where monetization could be advertising or consumers subscribing to content. You can perform multiple different monetization capabilities today. Viewership could be across the globe. This is the chain: starting from the camera, producing content, preparing the content, distributing the content, and eventually getting

it to viewers.

I want to give you a retail analogy. Like retail, there is a production phase where you produce goods, take them to a warehouse for packaging, send them through trucks and logistics to retailers, and eventually consumers buy them either directly or through e-commerce mechanisms. Very similar things are happening in the content business as well. If I am a content creator, the first thing I do is produce content. Exactly like a factory, there is a content factory where hundreds or thousands of hours of content are created. Once it is created, it goes through a preparation stage where, like a warehouse, you are preparing this media, and eventually, using logistics, you are sending it to the platforms that deliver it to viewers.

AI is transforming every industry, and the media industry is no different. We are going to understand this transformation across two spectrums. One is from the perspective of a content creator, a studio, or an OTT operator looking for operating leverage through AI—specifically agentic workflows—or market expansion where AI can accelerate reach and audiences. These are the two real levers we talk about when leveraging AI.

The first step is from an operating leverage standpoint: agentic workflows. If you go back to the production of consumer goods, the biggest change in the retail industry is the evolution of physical AI, like robots making factories autonomous so that the cost of production comes down. A very similar scenario is happening in the production space of the content business, which is being disrupted by Gen AI. We are starting to see high-quality production systems starting to evolve with tools like Runway and Sora. In the next few years, you are going to see a fundamentally shift. The cost of production is going to deflate.

From an economic standpoint, if the cost of content production deflates, the economics of the business must be rethought for every part of the value chain. We are starting to see a new set of AI-first studios evolving. These studios use AI as a tool but bring in the creative storytelling capabilities of humans to create compelling stories. We are seeing this across the globe as a major shift.

I feel live sports might be the only non-AI content to survive in the long run. Live sports are like gladiator sports; we love to see humans competing against each other. That is something we will continue to see given human emotions. But for almost all other content, it will eventually be hard for humans to delineate whether it was AI-generated or human-generated.

Once you create the content, the next important piece is the media preparation stage. This stage has a lot of human cost today. People must take the content, encode it in different formats, and create metadata. All the data you see on an OTT platform screen—the artwork, description, storyline, program guide ratings, actors, episodic information, content advisories, audio tracks, subtitles, creators, and PG ratings—forms the metadata.

Think about an environment where hundreds of thousands of hours of content are created for 40 countries, meaning you need multiple languages and must comply with different regulations. This is the heart of the metadata explosion. Today, this is a human cost. Lots of humans in this content factory are manually creating metadata or looking at compliance. If you are a sports or news company creating hours of unique content every day, you become inundated with the work. There

is clearly a headcount crisis in many media companies because of the volume.

When we talk to customers, we learn that for every $1 they spend on technology, there is $2 to $4 of human toil. This is a big limiting factor for video businesses because everyone needs "eyes and ears" to control and manage the system. This is where agents come to the rescue. The chore of content preparation and globalization will start to become agentic.

To explain what an agent is, take a coffee machine. A standard coffee machine is a non-agent; you press a button for an espresso or a latte, and it follows a pre-programmed sequence. This is just basic automation. An AI-enabled agentic coffee machine would monitor you and, based on your facial expressions and the time of day, predict that you might need an espresso. It proactively anticipates your needs and acts on them.

As AI technology becomes better, we are moving toward true AI agency, or the AGI world. In a true agency, the machine learns and decides by itself in ways creators do not anticipate. The machine might say, "I have invented a new drink for you that I think you will like." We are seeing this transformation starting to happen in the media business.

In the warehouse of media packaging, we are seeing an evolution of digital agents. These are AI-enabled software systems that understand video, audio, and storylines to extract and create metadata automatically. Machines can create story summaries without a human in the middle. These agents are replacing chore work and solving the problem of "messy media" which has previously slowed down scaling.

Imagine the daily experience of a scheduler for a TV channel. In an agentic world, between 2:00 AM and 5:00 AM while you are sleeping, the system is building schedules automatically. It understands your taste, the genre of the content, business needs, and yesterday's viewership analytics. It might even look at social signals—like noting it was a certain actor's birthday—and schedule accordingly. At 7:00 AM, the system provides a summary of the schedule and asks the human operator to make final decisions. A task that currently takes 8 hours could be done in 10 minutes. This allows a single person to manage dozens of channels, meeting the global demand for more content.

Once the content is ready, it must be delivered across multiple platforms globally. This is a logistics problem. When a studio transacts with an OTT platform today, it involves a lot of manual human interaction and emails regarding Diwali shows or other seasonal planning. We see a future where two agents talk to each other. An agent representing a platform can talk to hundreds of content creators to aggregate the best shows for a holiday and prepare advertising and promotion.

This leads to inter-company agent transactions. Agents within different companies will transact, negotiate commercials, and handle legal documents like a marketplace. This will be the core future of content distribution.

Regarding delivery to consumers, we currently discover content through websites or apps. We are moving toward agentic discovery. Most of us are already moving from search engines to conversational bots like ChatGPT or Gemini. These bots understand our nuances and preferences. Eventually, we will have a personalized conversational bot that understands our watching behavior across platforms like Netflix, Prime, and YouTube. It will understand our mood and emotions better

than anything else. Discovery will shift to these entirely different surfaces.

The whole value chain is changing. AI can also help extend reach and create entirely new sets of content. Today, in sports or news, we capture a rich reality with 20 to 100 cameras, yet we often see a homogenized, bottlenecked version because humans cannot scale the switching and production. With AI, we can transform this.

A single football match can be turned into many different stories. You could have a version for analytical fans, tactical fans, or home fans with different camera switches, commentary, and graphics specific to that persona. For a Gen Z audience, you could have a vertical screen version with social chat and multiple simultaneous camera positions. This native storytelling across different formats was previously too expensive but is now possible with AI.

We also see potential in micro-dramas. Production houses are creating these from scratch, but we can also repurpose existing content from the last 50 years. Imagine taking a 3.5-hour movie like Sholay and converting it into 10 episodes of 3-minute micro-dramas. This allows you to bring old stories to a new world through AI-driven transformation.

The technology to solve these problems is complex because it involves audio, video, and text. In video businesses, we need to understand music, background audio, and speech. We need to understand characters, action, and emotions—like identifying a dunk or a yellow card in a game. While LLMs are great for text, we need Vision Language Models (VLMs) to act as the "eyes and ears."

Current VLMs are still expensive and slow. The next step is "World Models," which predict the next moment in a real-world environment. This started with autonomous cars but is a huge use case for video. In sports, the system can understand the physics of the ground and air drag to predict a ball's trajectory. This will lead to truly immersive video. You could watch a game from the shoes of a player or the gloves of a goalkeeper. Immersive video is dependent on how world models evolve, and this will dramatically change the experience for audiences.

Because video requires low latency and high performance, we will see hundreds of custom audio-video models. Some will be specific to soccer or basketball, while others might focus only on modeling an ocean scene. This goes beyond what standard LLMs can do.

In summary, AI transformation is happening across the entire content value chain. This is a foundational transformation that happens once in a lifetime. It will change who creates content, how we consume it, and the underlying economics.

Operator: There is a QR code on the page for you to leave feedback. I will start by voicing some of the questions from the chat. Please note we are strictly focused on AI in media and will not take company-specific questions.

The first question is from Pratabi Dutta. Given that we have strict SLAs of 99.9999% with clients, agentic AI can be prone to hallucination and be unpredictable. What safeguards are in place to mitigate such risks?

Management: To make systems deliver value without hallucinating, companies must put a lot of guardrails in place. You build "eval" systems to bring determinism to an indeterministic problem. When dealing with hundreds of thousands of hours of content, predictability is the most important thing. Demos and PoCs are easy; the hardest part is getting them into production and making them work realistically across all scenarios. That is the core of engineering for AI solutions.

Shankar Narayanan S: Regarding the prep work, are we seeing any improvements in the speed of work due to AI? Does it in turn reduce pricing because of AI-led deflation?

Management: This is primarily a human cost reduction venture. The fundamental idea is that the human element of tasks—like speech-to-text for subtitling, dubbing, or artwork and promo creation—is changing. This is an incremental TAM (Total Addressable Market) because it addresses human costs that previously limited expansion. We see this as enabling newer business possibilities.

Shankar Narayanan S: Can AI become a lever for getting production and pre-production workflows to be hosted in the cloud?

Management: Absolutely. One trend we see is that AI is driving a progression toward scalable cloud platforms. Most facilities do not have sufficient on-premise access to the GPUs and servers needed for AI. The acceleration of AI workloads is a big driver for moving from on-premise to the cloud.

Anmol - DAM Capital: Have we seen an increase in OTT or cable-based content because of AI? Are there any real-life examples where costs have been reduced?

Management: It is currently happening in bits and pieces; it is not yet a broad factory approach. We are seeing content creators find that VFX costs are coming down. They are not changing the whole movie yet, but parts of the story now use AI for background creation or specific locations. It is early days, but in the world of micro-dramas, we see a lot more activity. The direction is clear: there will be a tsunami of AI-created content.

Ayush Shah: Are there any evidence-based cost savings for production and preparation, especially given the higher token costs?

Management: Token cost is only a small part of the cost of AI since this involves video, which uses custom models and GPU costs. While there are GPU expenses, humans are very inefficient at some of these tasks. The cost of a GPU is not comparable to the human cost required for the same volume of work. Most of our customers are looking at this for expansion and new revenue possibilities rather than just cost savings.

Rohan Nagpal - Helios: Are there any network effects that can come from agent-to-agent communication? Is there any benefit to being a first mover?

Management: If you can connect two parts of an ecosystem through agentic infrastructure, the network effect is extremely high. While internal agents enhance productivity and operating leverage, agents that coordinate across different entities have a dramatic multiplier effect. Every industry is currently going through the process of defining inter-agent communication, protocols, and standardization. The biggest value lies in the time compression of these transactions. Negotiating legal agreements between partners can be brought down by orders of magnitude in

time and effort. Most capabilities will be "agentified," while humans make the final key decisions.

Bharat Gulati – Dalal & Broacha: Does AI help strengthen the business through deeper engagement with clients? Does this create a threat to horizontal SaaS or hyperscalers?

Management: From a vertical standpoint, like the media business, we deal with mission-critical software. Context is the most important moat for any business. While reasoning might become commoditized, whoever owns the enterprise, technology, and operating context will drive the workflows of the future. Whoever owns that context will be the primary driver for these businesses.

Sharad Goenka: With the rapid adoption of agentic AI across media, do you expect it to materially reduce operating costs for your customers? If so, could this create pricing pressure?

Management: I see the Jevons paradox playing out. If you have more automation and capability, civilization tends to do more things rather than fewer. Customers are asking how they can use this to expand their revenue opportunities and capabilities. While the cost per individual job might deflate, the multiplier effect of the volume of jobs people want to do looks expansionary for the industry.

Chintan – Girik Capital: How do you see industry commercials between vendors and clients changing as the industry moves toward agentic AI deployment?

Management: It is very early, but people are starting to explore outcome-driven changes. We see this in other industries, like call centers, where pricing is based on a per-transaction or per-result basis. The core of systems will likely remain based on the existing value equation, but everything related to human cost leverage or new revenue sharing from expansion will likely be driven by outcomes. How pricing is determined linear to those outcomes is still playing out in the market.

Operator: We are at the end of the question queue. If you have any other questions, please feel free to write to us at IR@amagi.com. Thank you all for joining.

Management: Thank you, everyone. Please share your feedback on what we can do better or new topics you would like to discuss. We are fortunate to have a ringside view of this worldwide and are happy to bring in experts to help you understand this better. Thank you very much.