

What's the Difference Between the Enterprise and Consumer Metaverses?

The enterprise metaverse is a Fortnite but for digital twins. Presented is a perspective for developers on the challenges—and opportunities—ahead.

The “enterprise metaverse”—the metaverse for commercial, industrial infrastructure, and manufacturing applications—will be far removed from the social, gaming, and entertainment ecosystem that many of us regard today as the “consumer metaverse,” where entities like Fortnite and Roblox have staked an early claim.

We intuitively understand the applications for a consumer metaverse, where the core mechanics are focused on game play and socialization. But what are the applications for an enterprise metaverse, and what are they good for? Today we see three main focal points for enterprise metaverse applications emerging from our engagement with customers:

Synthetic Data Generation (Open-Loop Simulation)

The idea here is to create an environmental context where virtual sensors are placed either in the environment itself or on systems that inhabit it. The goal is to collect data—typically, perception data—for the purpose of training machine-learning (ML) models, or other forms of autonomous software.

Is it worth double-clicking on what we mean by an au-


tonomous system? It's a mechanical system that knows how to respond dynamically to its context. A robot or other mechanical system that's *not* autonomous will do exactly what it's programmed to do. What makes it autonomous or semi-autonomous is when it can make sense of its dynamic environment and then act accordingly—and independently.

But human experience has taught us that it can take a long time to learn how to operate successfully within a new context! The same is true for machines. The way to accelerate that learning is by exposing autonomous systems to more diverse information. It's not enough simply to generate *more* data. It's important that the data itself represents the diversity of situations that you want your autonomous system to respond to.

Collecting this data in the real world is often very expensive, time-consuming, sometimes dangerous—and in some cases, virtually impossible. It's not possible to order up edge cases on demand, yet they're exactly what autonomous systems need to be able to recognize and handle to operate safely and robustly.



Falcon, Duality's Digital Twin Simulator, utilizes digital-twin counterparts based off of real-life appearance, physics, and behaviors.



To facilitate machine-learning training, hyper-realistic synthetic data is necessary. System twins can be modeled with a suite of sensors, including depth, segmentation, and LiDAR.

Thankfully, when driving around town in our vehicles in the real world, traffic accidents seldom happen. Yet that's exactly what you'd want a self-driving car to experience and learn from. It just doesn't happen very often.

So how do you go and collect this data very intentionally? The enterprise metaverse holds this promise: The ability to generate very accurate, scalable, and intentional data for ML models.

Validation of Autonomous Systems and Processes Automation (Closed-Loop Simulation)

We can also think of the enterprise metaverse in terms of validation and optimization. This differs from Open-Loop Simulation/Synthetic Data Generation in that you don't merely have a virtual context in which a system is operating; you also have actuation of those systems. These systems can run their external software brain based on input from virtual sensors and then determine how they should respond. And you can test this in a real-time closed-loop simulation.

Thus, you can run many thousands of tests before you ever deploy a system like this in the real world. Whereas with a self-driving car, a change applied to even just one line of code could make the software—and vehicle—behave unpredictably or even crash (think of how common software bugs are...). We obviously can't afford for this to happen when we introduce autonomous systems into real-world operating environments.

How do we ensure that these systems are robust and validated? It's not realistically possible every time you change a line of code to go drive a thousand miles for validation. The same applies in a construction site or a warehouse. Every time you change a line of code in your mobile robot's logic, if that robot is deployed at 40 customer sites, you can't realistically test that robot at all 40 sites before you release the software. So, how do you actually validate these systems effectively?

Until recently, we've been thinking in terms of single systems. But now imagine that in a warehouse, or an airport, we have a broad collection of systems that are autonomous, semi-autonomous, or human operated, all working together in close orchestration. At this point, it becomes increasingly

challenging to think about how to go about optimizing a workflow like that from an operations perspective.

The enterprise metaverse holds this promise. For the first time, we're allowed a fully instrumented environment where a complex set of systems can interact and orchestrate, and subsequently analyze all of the output data in a closed loop to figure out what and where to optimize. Then, we can run A/B tests or myriad configurations in parallel to find the optimal workflow.

B2B2C Interaction

In some cases, we want to expose these digital twins to end customers or consumers. A simple example: a shaving blade. How do you really show off the features of that blade? Until recently, manufacturers have invested as much as 4X more money in the packaging than the shaving blade itself, all to influence the purchase decision. But how might you influence consumers' purchasing decisions when over 90% of those purchases are made online?

One way is to allow them to interact directly with a digital twin of your product, or your service, so that they can more fully understand and appreciate the value proposition. And this is where the lines between the enterprise metaverse and the consumer metaverse begin to blur.

Digital Twin Challenges

The three dynamics discussed above are what we've seen most frequently with customers, and all of them can drive significant business value. Some shift the bottom line, some the top line, but these are firmly grounded use cases addressing significant pain points for companies today. In other words, the metaverse future is already here.

It's a paramount concern to ensure that these autonomous "smart" systems aren't suddenly going to become "dumb" when you simply change a couple of lines of code. It's also important to allow your customers to interact directly with these digital twins so as to get a much better understanding of them.

Today there are online product configurators for almost every new automobile, a simple example of a digital twin. But it can cost half a million to a million dollars and six



Techniques like photogrammetry allow for digital-twin acquisition that's based on real-world data.

months working with a consulting agency to develop one of these configurators. How can we practically do this for an \$8 to \$12 product?

These are among the many metaverse challenges that enterprises face today. They're not solutions in search of problems, but rather acute problems that enterprises currently grapple with, and frankly don't have great solutions to overcome them.

Enterprise vs. Consumer Metaverse

With the applications now established, what are the key differences between the consumer metaverse and the enterprise metaverse? There are three main differences to consider.

Who is in the metaverse? In the consumer metaverse you'll find avatars—avatars that often represent the users who are in that metaverse with their friends, performing activities, playing games, etc.

The enterprise metaverse, on the other hand, consists of digital twins. These are digital counterparts that accurately capture the appearance, physics, and behavior of physical systems, items, and spaces making up the actors and context of enterprise workflows.

If you have a digital storefront, the store is an **environment twin**. The products on the shelves are **item twins**. A re-shelving robot or a frictionless checkout counter would be **system twins**. The number one difference from the consumer metaverse is that the enterprise metaverse isn't predominantly a metaverse of humanoid avatars. It's a metaverse of digital twins.

Permeability of the data between the real and the virtual: In a consumer metaverse, what matters most is that the metaverse itself is engaging, whether it follows real physics or takes creative liberties (lighting, shading, physics, etc.). Fortnite, Roblox, and Minecraft **embrace** a very high degree of creative expression and stylization. They're not meant to resemble the real world, and this is by design.

An enterprise metaverse, on the other hand, **must be precise**, mainly because it must generate data intended to be permeable with the physical world. For a simple example, envision the digital twin of a dump truck. Whether it's driving through the terrain or dumping its load, the speed and

physics must behave the same way in both the metaverse and the real world. Otherwise, the data produced by the metaverse just isn't meaningful for machine learning or validation.

Now envision a consumer shopping interaction. If a digital twin of your product behaves differently than it would in the real world—perhaps it's a garment that responds to lighting in a different way as it moves—then the experience is largely useless. The consumer can't make an informed purchasing decision based on that interaction because the metaverse product isn't representative of the real-world product.

Data generated in the enterprise metaverse must therefore be permeable, which is a two-way street. The way these digital twins are created and tuned to work correctly is typically based on data that's harvested from the real world. The enterprise metaverse is created and tuned using real-world data, and the behaviors in *that* world also produce data that is meaningful in the real world.

What is the metaverse for? The enterprise metaverse isn't an interaction-centric metaverse. It's a data-centric metaverse. The very reason for operating the metaverse is to produce data of some kind. This could be learning data for ML, or validation data for an autonomous system.

Even when consumers are interacting with the enterprise metaverse, they're creating data to evaluate how that product will work for them in the real world. The value proposition here isn't entertainment, it's functional.

Therefore, the goal of the enterprise metaverse is really data generation. In that sense, two things become very important: The entire enterprise metaverse must be heavily instrumented for efficient, scalable data generation. And, because the data doesn't live on its own, it must have a clean way to integrate and flow into and coexist with other enterprise data and analytics pipelines.

Easy, Right?

Some major challenges emerge when attempting to make all of this work elegantly at scale.

Achieving data permeability

On the face of it, it sounds simple: In the enterprise meta-

verse, things should behave just like in the real world. But think of the complexity and diversity and variety of what can happen in the real world—it's essentially unbounded. One has to be very focused and disciplined about what data permeability really means in a given context.

In what domain does the metaverse need to operate? Even after that's defined, it's essential to ensure that the digital twins behave accurately, and virtual sensors produce data at the desired speed and fidelity. What's typically required here is *lots* of deep domain knowledge—knowledge of simulation, physical systems, and mechanical systems—to truly make all of this work. A diverse range of skills, talents, and expertise in the team, vendors, and partners is critical.

Customer learning

Customers are very comfortable today working with data and analytics, and harnessing them to impact how key decisions are made. It's how any successful business operates today.

AI and ML are just starting to achieve that level of mainstream understanding and adoption. Oftentimes, even when we don't think we're using ML, it's there under the hood and it's solving practical challenges in seamless ways within current workflows.

When it comes to digital twins, companies today are just beginning that journey. The tools, pipeline, and plumbing doesn't really exist. The early adopters (brave souls!) will forge the way forward for everyone else. There's a significant learning curve for companies just starting to adopt digital twins in their workflows, and it's best to acknowledge and plan for that.

Standardization and decentralization

The third challenge is perhaps the trickiest to solve, with major implications for metaverse technology and commerce in the future.

Today, when we talk about the consumer metaverse, you don't think of a lonely place to occupy by yourself. You'll be with family, friends, colleagues, sometimes strangers. The metaverse is enriched by these interactions with others.

From the enterprise view, if companies want to bring together assets from different teams, other companies—perhaps from their own supply chain—can be incredibly difficult because they understandably control their assets very carefully. Bringing physical assets together become massive logistical challenges—legal friction, logistical friction, protections that are there for a reason.

But different companies, different teams—sometimes different teams within the same company—can learn a lot by bringing their assets together in a shared context. The enterprise metaverse has the promise to fulfill this potential. It represents a brand new way for companies to think, work, and operate if the friction to collaboration and developing shared contexts can be removed.

This, however, requires that the metaverse itself doesn't wind up being a walled garden for each company, or a set of walled gardens from different vendors. Digital twins in the enterprise metaverse must be expressed in a standardized and decentralized way so that they can be exchanged with others.

Unfortunately, on the consumer side, the trend is very much aligned toward walled-garden metaverses where it's not possible to exchange assets between them. Consumer metaverses *want* that consumer lock-in.

On the enterprise metaverse side, there's significant value for participants to not be locked in. Our hope is that the economic value the participants can potentially gain from having a common metaverse construct—a standardized digital twin format—will result in a metaverse without walls. The vector we want to push for is standardization in terms of formats and interoperability, and decentralization.

The enterprise metaverse—the Fortnite of digital twins, if you will—presents significant challenges, but also a significant opportunity. There's real value to customers participating in standardized, interoperable, and decentralized platforms.

There will be plenty of opportunity to generate revenue in the enterprise metaverse, but it won't and shouldn't be achieved at the expense of locking in customers into walled gardens. This only hurts the ecosystem and reduces the overall value that enterprises can gain from the metaverse.