

# TechTalk

We are proud to be recognised as the  
**"Overall Partner of the Year  
Award - APAC"**

by Candefero at #CanalysForums2022,  
Singapore and bag this prestigious  
award.



**'BEST SYSTEM INTEGRATOR'**

at Star Nite Awards 2022  
by VARINDIA.



**Anoop Pai Dhungat**  
Chairman & Managing Director

**MD SPEAKS**

Dear Readers,

*I am pleased to inform you that Galaxy has been recognised as the Canalys Overall Channel Partner for the Asia Pacific region. This is indeed a proud moment for all of us at Galaxy and reinforces that we are on the right track. On behalf of all at Galaxy, I thank all our stakeholders because without their support this would not have been possible.*

*It is our constant endeavour to bring the best-of-breed technologies to our customers after enabling our resources in the same. We have now included a few more offerings in our portfolio that can provide you backup solutions as a managed service or even as software as a service. These solutions are extremely cost effective and reliable. Please call our specialists who will be able to explain these solutions in detail and customise and demonstrate the applicability to your organisation.*

*As we get to the end of 2022, I wish you very happy festivities and a wonderful beginning to the new year in advance.*

*Thank you again for your continued support.*

Happy Reading

*AP Dhungat*



# Future Is Now

## Space hotel scheduled to open in 2025

Waking up in a chic hotel room with a view of the solar system could be the future of travel. The US-based company has revealed new information and concepts for its space hotel idea, designs for which have been orbiting since 2019.

Originally premiered by Californian company the Gateway Foundation -- and then called the Von Braun Station -- this futuristic concept consists of several modules connected by elevator shafts that make up a rotating wheel orbiting the Earth.

Orbital Assembly is now aiming to launch not one but two space stations with tourist accommodation: Voyager Station, the renamed original design, is now scheduled to accommodate 400 people and to open in 2027, while new concept Pioneer Station, housing 28 people, could be operational in just three years.

The goal, says Orbital Assembly, is to run a space "business park" home to offices as well as tourists. Space tourism seems closer than ever before -- over the past year, billionaire Virgin founder Richard Branson blasted into suborbital space with his company Virgin Galactic, while Star Trek actor William Shatner became the oldest person in space thanks to a jaunt with Blue Origin.

But there's still a pretty unbelievable price point attached to any space trip, which makes it hard for many of us to actually envisage spending our annual leave out of this world. Tim Alatorre, Orbital Assembly's chief operating officer, thinks this barrier will lift as space tourism takes off.

"The goal has always been to make it possible for large amounts of people to live, work and thrive in space," Alatorre told CNN Travel in a new interview. In a 2019 interview with CNN Travel, Alatorre explained the physics of Voyager Station as working like a spinning bucket of water.

"The station rotates, pushing the contents of the station out to the perimeter of the station, much in the way that you can spin a bucket of water -- the water pushes out into the bucket and stays in place," he said.

Near the center of the station there would be no artificial gravity, but as you move down the outside of the station, the feeling of gravity increases.

The physics haven't changed, said Alatorre more recently. But, he explained, as Pioneer Station will be smaller, its gravity level would be different. There will still be what he calls the "comforts" of artificial gravity, like showers, the ability to eat and drink sitting down -- but the spaces with less gravity will allow for even more fun, space quirks.



<https://cnn.it/3Uv3E45>

## What is a data lake?

A data lake is a storage repository that holds a vast amount of raw data in its native format until it is needed for analytics applications. While a traditional data warehouse stores data in hierarchical dimensions and tables, a data lake uses a flat architecture to store data, primarily in files or object storage. That gives users more flexibility on data management, storage and usage.

Data lakes are often associated with Hadoop systems. In deployments based on the distributed processing framework, data is loaded into the Hadoop Distributed File System (HDFS) and resides on the different computer nodes in a Hadoop cluster. Increasingly, though, data lakes are being built on cloud object storage services instead of Hadoop. Some NoSQL databases are also used as data lake platforms.

### Why do organizations use data lakes?

Data lakes commonly store sets of big data that can include a combination of structured, unstructured and semi structured data. Such environments aren't a good fit for the relational databases that most data warehouses are built on.

Relational systems require a rigid schema for data, which typically limits them to storing structured transaction data. Data lakes support various schemas and don't require any to be defined upfront. That enables them to handle different types of data in separate formats.

As a result, data lakes are a key data architecture component in many organizations. Companies primarily use them as a platform for big data analytics and other data science applications requiring large volumes of data and involving advanced analytics techniques, such as data mining, predictive modeling and machine learning.

### Data lake architecture

Many technologies can be used in data lakes, and organizations can combine them in different ways. That means the architecture of a data lake often varies from organization to organization. For example, one company might deploy Hadoop with the Spark processing engine and HBase, a NoSQL database that runs on top of HDFS. Another might run Spark against data stored in Amazon Simple Storage Service (S3). A third might choose other technologies.

Also, not all data lakes store raw data only. Some data sets may be filtered and processed for analysis when they're ingested. If so, the data lake architecture must enable that and include sufficient storage capacity for prepared data. Many data lakes also include analytics sandboxes, dedicated

storage spaces that individual data scientists can use to work with data.

### However, three main architectural principles distinguish data lakes from conventional data repositories:

- ▶ No data needs to be turned away. Everything collected from source systems can be loaded and retained in a data lake if desired.
- ▶ Data can be stored in an untransformed or nearly untransformed state, as it was received from the source system.
- ▶ That data is later transformed and fit into a schema as needed based on specific analytics requirements, an approach known as schema-on-read.

### What are the benefits of a data lake?

Data lakes provide a foundation for data science and advanced analytics applications. By doing so, they help enable organizations to manage business operations more effectively and identify business trends and opportunities.

For example, a company can use predictive models on customer buying behaviour to improve its online advertising and marketing campaigns. Analytics in a data lake can also aid in risk management, fraud detection, equipment maintenance and other business functions.

Like data warehouses, data lakes also help break down data silos by combining data sets from different systems in a single repository. That gives data science teams a complete view of available data and simplifies the process of finding relevant data and preparing it for analytics uses. It can also help reduce IT and data management costs by eliminating duplicate data platforms in an organization.

A data lake also offers other benefits, including the following:

- ▶ It enables data scientists and other users to create data models, analytics applications and queries on the fly.
- ▶ Data lakes are relatively inexpensive to implement because Hadoop, Spark and many other technologies used to build them are open source and can be installed on low-cost hardware.
- ▶ Labour-intensive schema design and data cleansing, transformation and preparation can be deferred until after a clear business need for the data is identified.
- ▶ Various analytics methods can be used in data lake environments, including predictive modeling, machine learning, statistical analysis, text mining, real-time analytics and SQL querying.



## 5 Ways Data Protection for Kubernetes Is Different

The differences between Kubernetes data protection and traditional offerings go beyond data protection, highlighting key strengths of cloud native computing.

What are the differences between Kubernetes data protection and the more traditional data protection offerings, even though such products deal with virtualization and cloud scenarios as well as traditional on-premises data protection?

Here are five salient differences between such products — differences that go beyond data protection, highlighting some of the fundamental strengths of cloud native computing generally.

### **Data Protection Centering on Metadata vs. Data Protection Centering on Data**

At the heart of data protection lies backup and restore functionality. There is more to data protection than these two capabilities, but without backup and restore, there is no protection whatsoever.

In traditional environments, which for the purposes of this article include virtualization and cloud as well as various on-premises environments, backup and restore are focused on persistent data and the storage that contains it.

Kubernetes data protection, in contrast, focuses on metadata as well as the underlying data.

Kubernetes is essentially a declarative, configuration-based container orchestration platform. Providing data protection to those configurations and other metadata, including resource definitions, Helm charts and other files is central to the Kubernetes data protection challenge.

### **Dynamic Policies for Auto-Discovered Applications vs. Static Policies for Predefined Applications**

Setting up a traditional data protection application consists of establishing a set of backup and restore policies that apply to the various resources the organization wishes to protect. Such policies generally center on snapshots and backup schedules.

In Kubernetes, applications and their microservices

components are inherently ephemeral — scaling up and down at a moment's notice, occasionally appearing and disappearing altogether.

A Kubernetes data protection product like must therefore auto-discover applications on the fly to know what data and metadata to protect. The policies that drive such protection must correspondingly be dynamic as well.

Dynamic policies exist at an abstraction layer above static ones, and the underlying technology must interpret them in real time in order to apply them properly in each situation.

### **Dynamic, Policy-Driven Automation vs. Static, Manually Configured Automation**

When the Kubernetes environment interprets and applies policies, what it's really doing is automating workflows that those policies specify.

Traditional data protection technologies also feature policy-driven automation, but those automations are as static as the policies themselves.

Kubernetes thus requires a rethink of what automation means — instead of a simple flowchart of “do this, make a decision, and do that” logic, cloud native automation is inherently dynamic, with logic that might change from moment to moment.

This revamped notion of automation applies to data protection as well as other Kubernetes automation scenarios.

### **Application-Specific vs. Volume-Specific Data Protection**

Because traditional data protection centers on data and storage, operators logically focus on backing up and restoring databases and storage volumes.

Volumes, in fact, are the common denominator for all traditional data protection, since backing up and restoring them means backing up and restoring anything stored or installed on them, including databases, files or application components.

Kubernetes, in contrast, maintains a comprehensive, declarative abstraction of the entire persistence tier.



# Special Focus

Kubernetes applications are fundamentally stateless, given the ephemerality of their components and the stringent performance requirements that apply to such applications.

Nevertheless, Kubernetes applications must typically maintain state without affecting these core characteristics. The platform overcomes this challenge via configuration-based abstractions.

This abstraction-based state management means that Kubernetes data protection cannot take place at the volume layer. It must take place at the application layer instead because only applications know what data they need and when. Details about storage have been fully abstracted away.

## Application Recovery vs. Data Recovery

The most important principle of data protection is that your backups are only as good as your ability to recover from them.

While data recovery is most of the traditional data protection recovery story, Kubernetes recovery involves a complicated combination of the recovery of the data, resource and configuration components that make up a running application.

Such recovery involves the automated orchestration of several dynamic, policy-driven tasks – a tricky proposition given that the goal of such recovery isn't simply the avoidance of loss of data, but rather, the continuous execution of applications in production while minimizing any adverse impacts on the users of those applications.

Once you understand the full complexity of such automations, it becomes clear why Veeam acquired Kasten. Delivering on the full Kubernetes data protection value proposition, as the Kasten K10 data management platform does, is no simple task.

## The Intellyx Take

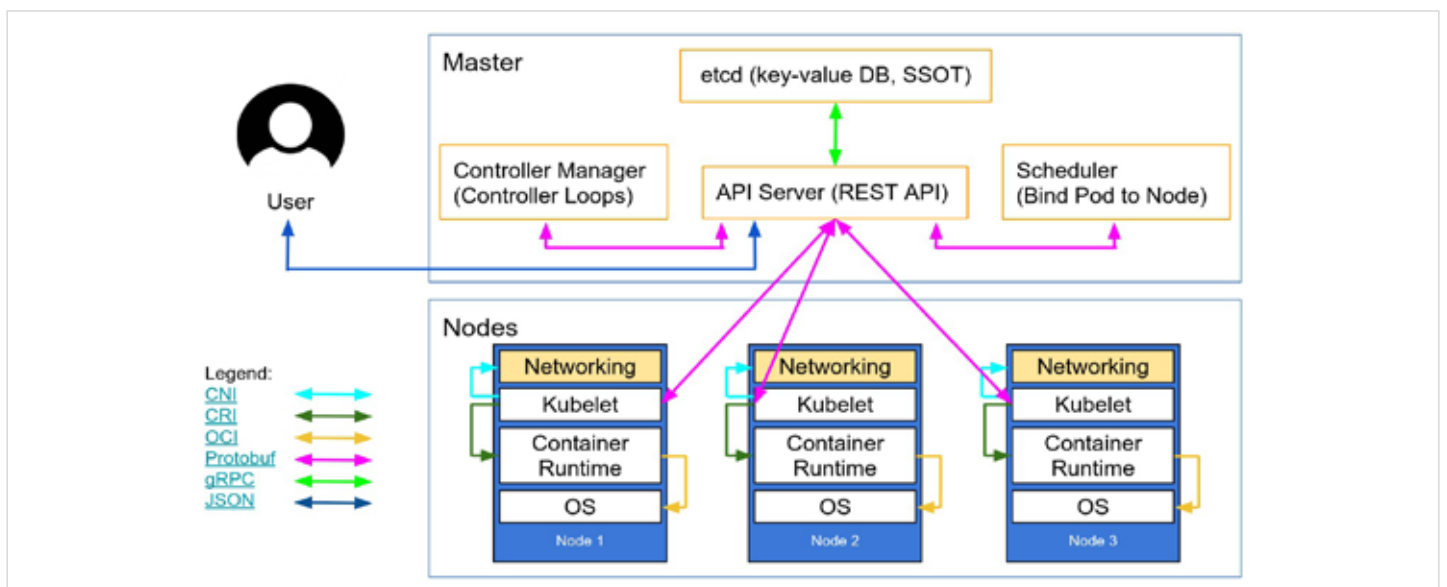
Highlighting the differences between traditional and Kubernetes data protection inevitably highlights the differences between traditional and cloud native computing.

Cloud native infrastructure requires a comprehensive declarative abstraction layer that abstracts storage and data, and enables stateless application behavior while managing state.

An important benefit of this cloud native approach is a clear distinction between control and data planes — the configurations that drive the behavior of applications vs. those that concern the movement of data.

Data protection must work at both layers: firstly, moving data as part of the backup and restore processes, and secondly, the broader data protection story that takes place as part of the automated orchestrations that drive application behavior.

At Galaxy, we have Certified Kubernetes Administrators who can assist you with the planning, migration, containerisation and implementation of your applications on the Kubernetes platform.





## Google launches its own text-to-video AI tool Imagen Video

Google answers Meta's "text-to-video AI, the "Make-a-Video," with its own, the Imagen Video. Researchers at Google Brain, the company's AI lab, introduced Imagen Video, an AI that can create video clips from text prompts. The second text-to-video AI comes six months after DALLE-2, a text-to-image generator from OpenAI, and merely a week after Meta announced its "Make-A-Video."

Google's Imagen Video can produce videos of 1,280x768 pixels resolution at 24 frames per second of not more than 5.3 seconds. The model takes a description and generates a 16-frame, 3-fps video having 24 x 48-pixel resolution. Then, the system upscales and "predicts" additional frames, producing a 720p video at 24 frames per second.

Google says Imagen Video has a "high degree of controllability" and world knowledge. "We find Imagen Video not only capable of generating videos of high fidelity, but also having a high degree of controllability and world knowledge, including the ability to generate diverse videos and text animations in various artistic styles and The Imagen Video was trained with an "internal dataset" of 14 million videos and 60 million still images, and the training data further contained another 400 million images from the LAION-400M open dataset.

The team at Imagen Video plans to join the researchers at Phenaki, another text-to-video AI from Google that can turn detailed text prompts into two-minute-plus videos, though with a lower quality.

The demos shared include a video of "Coffee pouring into a cup," "Wooden figurine surfing on a surfboard in space," "Balloon full of water exploding in extreme slow motion," and more.

<https://bit.ly/3P6cG6s>

## India IT spending to grow 2.6%, public cloud spending to grow 27%: Gartner

India IT spending is projected to grow 2.6% in 2023, according to a recent forecast by Gartner, Inc. Even with the looming tensions of global inflation and the weakening rupee, Indian businesses will continue to increase their spending in key segments of information technology next year.

"Inflation has not impacted enterprise spending on technology globally, and India is no exception to this trend," said Arup Roy, VP Analyst at Gartner. "The headwinds are in favor of technology as businesses realized how going digital can benefit them in the long run.

Depending upon the maturity level of the digital enterprise, the spending context may be different for different businesses, but overall technology spending will continue to be on the rise in 2023."

Gartner forecasts weakening demand for devices in 2023 as device upgrades stabilize (see Table 1). Spending on data center systems will also experience a decline as businesses opt for public cloud services. End-user spending on public cloud in India is forecast to grow 27% in 2023. Except for data center systems and devices, all the other segments of IT spending will experience growth in India in 2023. Despite software segment's projected growth of 13.7, it is still a sharp drop from 19.9% in 2021. Roy states several factors contributing to the decline.

"Economy softening, geopolitical conflicts, no respite in the supply chain constraint, and the ongoing talent shortage contribute to the same," he states. "As Indian organizations advance in their digital initiatives and maturity, they will need to elevate their cybersecurity focus or face severe business risk. Also, paucity of good quality digital talent and technology/management skills will continue to be some of the top challenges that CIOs in India face in 2023," said Roy. CIOs in India will need to be extremely creative about how they source talent, where they source from and how effectively they use those resources.

<https://bit.ly/3VC7nhF>

*All product names, logos, brands, trademarks, and registered trademarks are property of their respective owners.*