Turn your Database into a Service

A dbInsight white paper for Nutanix

Trigger

IT is under tremendous pressure to deliver new application services that add value and spend less time on infrastructure. The cloud offers the promise of simplifying IT operations and most enterprises are now actively embracing the cloud. However, until recently the full promise of simplification from the cloud required enterprises to adopt managed services that only run on the cloud provider's platform. Now enterprises are demanding freedom of choice.

Most enterprises are shying away from only using one type of cloud and instead will use multiple clouds to prevent cloud vendor lock-in. This is especially true with managing data – as collecting and analyzing it for business purposes is becoming more important for enterprises to remain competitive. Today, most enterprises have highly varied database portfolios, with the certainty that no single database platform will address all of their needs. Most enterprises are working with data that originates from multiple sources inside and outside the firewall. Yet, gaining the best of both worlds – harnessing the operational simplicity of the cloud and the freedom to choose the right database platform for the right use cases – has been virtually impossible until now. The challenge is that the operational simplicity promised by the cloud until now has only been available from specific cloud providers, forcing enterprises to make difficult choices.

Our Take

By the mid-2020s, Hybrid Cloud Computing will become the default choice for deploying new systems or replatforming existing systems. Enterprises demand the simplicity, flexibility, elimination of operational and management duplication, and agility of *the cloud control plane*, while keeping their options open on where to run the workloads. For databases, they want the cloud operational simplicity of a service without restricting their platform choices. There are multiple paths emerging for enterprises seeking to manage their data estates through cloud-style consolidated control planes, with or without having to rely on specific cloud platform vendors. Hyperconverged infrastructure has emerged as one of those options. It allows enterprises to use the hardware of choice by applying a software-defined control plane that manages each of these pieces as modular building blocks. For HCI to become practical for bringing cloud operational simplicity to databases, they require highly simplified single-click solutions addressing the core housekeeping tasks of operating databases, encompassing provisioning, copy data management, data protection, and management of patches and updates that work across all of their databases, regardless of where they are deployed.

Figure 1. Evolution of enterprise cloud adoption

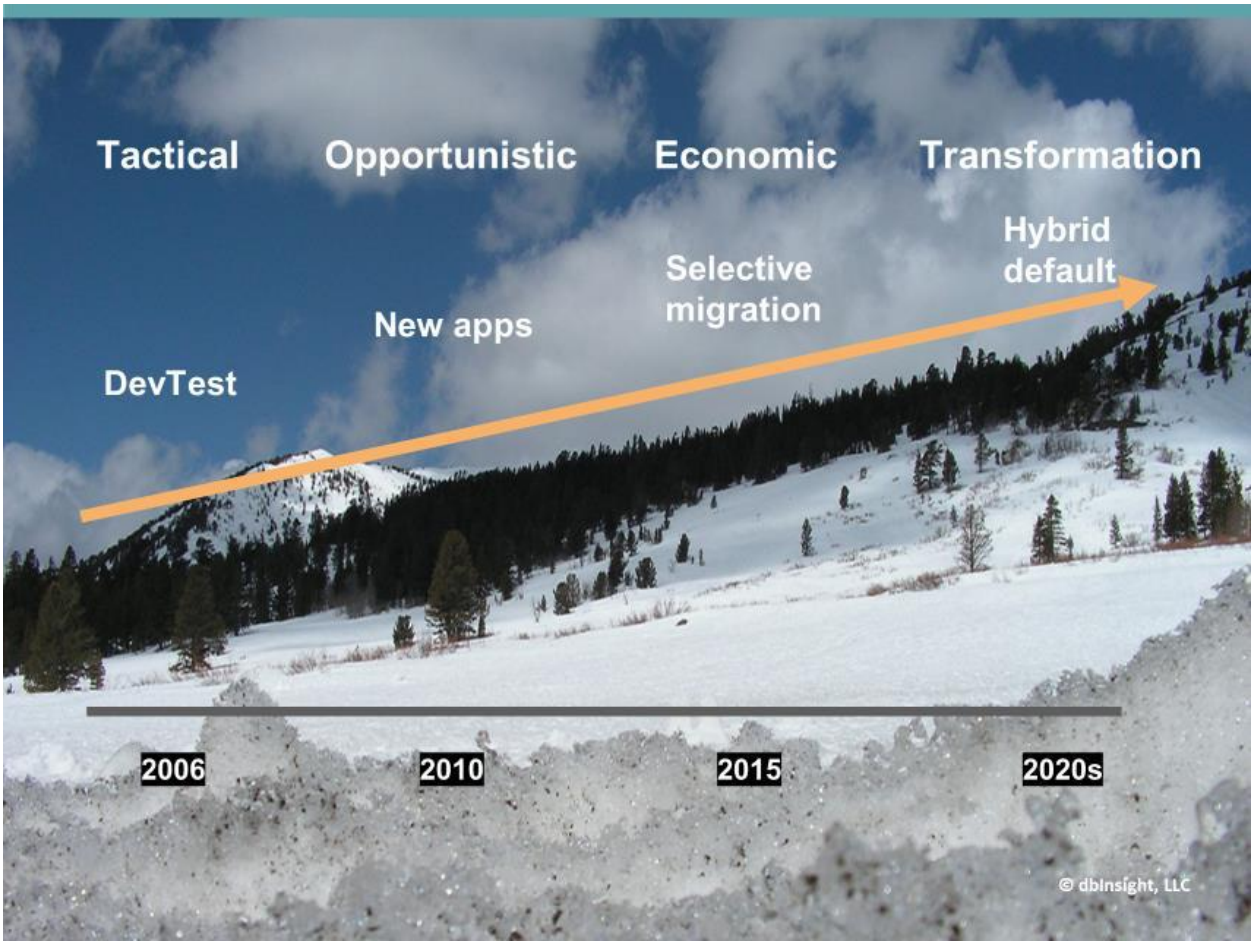

Source: dbInsight

The Hybrid Default

Cloud computing is already top of mind. According to a Flexera survey, 94% of respondents are using the cloud, of which almost all of them (84%) have a multi-cloud strategy. Cloud adoption has also evolved over the years, having begun with app developers with DevTest, followed by opportunistic development of new apps. Today, enterprises are actively embracing next-generation SaaS services that complement and update their enterprise applications portfolios in areas such as operational optimization, IoT and edge computing, and the emerging wave of AutoML services.

This represents a marked shift from traditional decision-making for deploying IT systems, where on-premises was considered the default option with the need to justify

the cloud alternative. By the mid-2020s, this process will flip. Enterprises won't necessarily move all of their systems to the cloud, but by the mid-2020s, they will have to justify the exceptions.

The drivers are:

- Demand for innovation – changes in the competitive landscape are driving the need for new applications supporting operational efficiency, more effective customer/trading partner engagement, more bulletproof security, and a generation of new insights that cannot easily be addressed with the backlog and management burden of legacy three-tier deployment.
- Expansion of the data estate – Apps that support innovation are increasingly reliant on nontraditional data types and sources originating outside existing enterprise transactional data sources – e.g., messaging, text, images, logs, sensors/IoT, social networks, and other data sources.
- Operational simplicity – IT organizations that spend the majority of time in maintenance (keeping the lights on) require new approaches that deliver the operational simplicity of cloud-based deployment.

The Data challenge

The data estate is growing more complex

Most already have hundreds of active production instances and multiple copies of each database. Copies are typically made to make data more accessible to more workgroups. However, they add significant management overhead in synchronizing multiple versions of the truth, and ensuring that they are consistently configured, patched, maintained, and secured.

For instance, lacking the ability to consistently enforce security policies governing what each class or role is allowed to see can result in leakage or exposure of PII or other confidential data to unauthorized users. This not only can lead to breaches, but also act as a brake to agility or responsiveness as IT or business users aim to mask or encrypt data after the fact.

And in most cases, it is not the same database. Most enterprises have “one of everything” in their portfolios as the result of line organizations or shadow IT making their own decisions and for enterprises that have undergone M&A have even more diversity thanks to differing platform standards. This complexity extends beyond the

datacenter to remote and branch offices, disaster recovery sites, and increasingly, edge computing.

It's also growing worse with the emergence of new use cases such as extending Customer 360 tracking, the dynamics of social networks, and building globally distributed transaction systems to satisfy emerging data sovereignty mandates. The increase in demand for new data platforms such as specialized columnar analytic databases, globally distributed transaction databases, NoSQL databases, graph databases, and data lakes also add to the growing data complexity.

In today's competitive environment, where organizations require real-time, reliable, common views of their customers and operations, having multiple sources of the truth is becoming a huge competitive liability – and exacerbating management bottlenecks.

Financial burdens of legacy architectures are preventing innovation

Legacy architectures impose financial burdens. It starts with the large upfront capital purchases that are required for small-scale deployments. The problem is compounded by the overhead of system maintenance. Performing mundane tasks becomes time-consuming. For instance, provisioning a database starts with a service ticket, then goes to the database administrator (DBA) who checks the database version and size of the resource. It then goes to the infrastructure team to configure compute and storage to create a cluster. It finally goes to the Dev/Test team for shakedown. This process can easily take hours at best, if not days or weeks at the worst.

The cloud opportunity

The cloud provides an opportunity to simplify database deployment and operation through Database-as-a-Service (DBaaS) offerings that automate provisioning, patching, updating, replication, and other housekeeping tasks. By choosing these services, enterprises gain the simplicity of the cloud environment.

But there are also several constraints to following the DBaaS route. First, by choosing these services, customers are actually making platform decisions because they are choosing a specific database/data platform with unique features that are only available from a particular cloud provider. Furthermore, while enterprises are embracing cloud provider DBaaS offerings, few mid/large enterprises will migrate all of their databases to the cloud.

In actuality, customers are demanding the operational simplicity and flexibility that a cloud control plane offers. Enterprises facing strategic deployment choices want to keep

their options open as they rationalize the diversity that exists (and will never disappear from) in their IT environments. And, in the data estate, they are demanding solutions that provide common control plane over their data silos regardless of where those databases run.

The role of hyperconverged infrastructure

HCI is a software-defined approach to consolidating and virtualizing IT infrastructure, from compute to networking and storage. It combines commodity datacenter hardware with locally attached storage to replace legacy infrastructure (consisting of separate servers, storage networks, and storage arrays) with virtualized, flexible building blocks. These building blocks lay the foundation for operating enterprise infrastructure as a private cloud.

By managing all components (compute, storage and networking) as singular infrastructure, HCI presents a compute environment that can be managed with the simplicity of the cloud. HCI enables flexibility, as many enterprises are already juggling running databases on-premises and in one – or more – clouds. It promotes flexibility and agility by eliminating the need to separately configure and tune each of the resources regardless of where they run, individually. Compared to traditional on-premise deployment infrastructure, HCI improves TCO by replacing costly legacy hardware (e.g., servers, storage networks) with off-the-shelf commodity infrastructure, and harnessing the scale-out power associated with cloud infrastructure.

Where Nutanix plays

Delivering a software-defined IaaS anywhere

Nutanix Hyperconverged Infrastructure (HCI) uses hyperconverged infrastructure to converge and virtualize the entire datacenter stack, including compute, storage, and storage networking. It provides software-defined infrastructure that delivers the benefits of cloud computing in the environment of your choice and in your control. Initially available for operation inside the datacenter, Nutanix plans to broaden the choice to also support the deployment of your organization's public cloud(s) choice. Nutanix HCI replaces complex legacy infrastructure and the complicated toolchains that are used to run it with an Infrastructure-as-a-Service (IaaS) environment that runs on commodity hardware. It offers dramatic operational agility and thus lower TCO by simplifying operations of various technology layers.

Its core components include:

- AOS, which provides a distributed data plane for either VMs or container-based applications that runs across a cluster of nodes delivering enterprise storage and virtualization services.
- Prism, which provides the distributed management plane that uses heuristics to simplify common workflows, eliminating the need for separate management solutions for servers, storage networks, storage, and virtualization.

Building on its core HCI platform, Nutanix has recently branched out to address the requirements of managing the application lifecycle, disaster recovery, object and file storage, IoT, and database management.

Nutanix Era, delivering the database- and cloud-independent DBaaS

Nutanix Era builds on the Nutanix HCI to deliver a DBaaS environment, delivering the best of both worlds: bringing the flexibility and operational simplicity of DBaaS without the restrictions of single-vendor solutions that are database, cloud provider, or location/deployment-specific. Currently supporting all major enterprise databases that are on-premises and/or on bare metal instances in the public cloud, Nutanix plans to extend Era to run across multiple clusters in major public cloud environments. The control plane will be the same for databases that are run anywhere, with disaster recovery and replication to span across on-premise, private, hybrid, and/or public cloud environments.

Era replaces the complex toolchains associated with provisioning, operating, and managing the lifecycle of enterprise databases through a web-scale architecture with a single control plane that scales transactional & analytical databases, linear performance scaling, and localized IO that offers ultra-low latency to database IO.

Nutanix Era delivers operational simplicity through a single click experience for common database management tasks currently encompassing:

- New database and database server provisioning based on best practices that are callable from a service catalog;
- Copy data management, that efficiently provides access to current data with a minimum of overhead;

- Data protection (backup/recovery) that builds in efficiency by leveraging snapshots and logs stored in Era, along with access to archival backups from native repositories of third party solutions; and
- Database patch management that are applied based on best practices.

Among the tasks covered are provisioning, cloning, patching, refreshing, and backups that offer convenience (with most operations performed in just a few minutes), a standardized environment (operating from the same versions of the software), and manageability and accountability (as all changes to the environment are tracked, making them audit-ready).

Era provides convenience with an integrated copy data management, snapshot, and data reduction capability that can create clones at any point in time, regardless of database size. Unlike most cloud-native databases, which replicate multiple physical copies, Era takes an optimized approach that minimizes its footprint. First, the clones do not consist of full copies, but instead comprise pointers to the most current version of the data. Secondly, Era also applies compression that reduces the physical volume of data by a factor of 5 – 6x. It also unifies patch management with a database-aware process that can minimize or eliminate downtime while keeping the system patched and in compliance with gold standards.

These features combined address a major operational pain point for organizations that otherwise would run multiple copies of the same databases, providing ready access to data, but at the cost of the waste and duplication. And thanks to the features reducing the footprint of database clones, Era generates significant savings in storage.

How Nutanix Era simplifies database operations

Era hides the complexity of database operations and provides common APIs, CLI, and GUI for multiple database engines. Here are a couple examples:

Era's Time Machine, the feature that handles copy management, backup, and restore functions, provides a good example. It creates a service layer for taking and recovering from backups through a combination of Nutanix platform snapshots and log backups that complements your existing backup solution; recent snapshots are managed by Era's Time Machine, while archival backups can be sent to the native repository of your existing solution.

You can specify the SLAs (e.g., retention period for logs) for routine database recovery operations when you register the database in Era. It then automates the process through:

- A self-contained registration process that discovers the logical layout of the database, maps the underlying physical disk resources, creates the protection domain and time machine entities for maintaining the SLA, and once the database is registered, taking the first snapshot using product-specific APIs.
- Automatically collecting logs, adhering to best practices approved by the database vendors.
- Use of full synthetic backups or snapshots that updates data collected from older backups with incremental updates that provides near instantaneous performance plus the ability to do point-in-time restores without the need of a full backup.
- The time machine capability captures essential data via snapshots (synthetic full backups) and transaction logs, preserving well-established database data protection practices while making them more efficient.

Another productivity feature is the built-in catalog of best practices for managing the lifecycle of the database, a feature that eliminates significant guesswork. Today, it offers pre-populated best practices for all databases supported by Era including Oracle, SQL Server, MySQL, and PostgreSQL. It provides the one stop for understanding how to best perform functions such as patching, and is extensible for your team to add their own best practices.

The complex handoff between DBAs, system operators (SysOps), and Dev/test are eliminated as DBAs and developers gain the one-click simplicity that cloud DBaaS services deliver. They can avoid all the handoffs, configuration, and testing operations associated with legacy on-premise deployment.

Takeaways

For most enterprises, support of a cloud control plane will drive deployment decisions in the 2020s. Enterprises are expecting the operational simplicity and agility that cloud computing environments deliver, while keeping their options open for supporting any database on any environment, on any infrastructure. With enterprises seeking to become data-driven, having this operational simplicity and flexibility in the database environment is becoming more important than ever. Enterprises want to transform their

databases into cloud-like services without restricting their platform or deployment choices.

DBaaS offers the promise of making database management simpler by providing a common control plane that automates the housekeeping tasks of provisioning, snapshot replication, high availability and failover, patching, and upgrading.

Nutanix meets this requirement by providing a software-defined platform that helps IT to offer simplicity of a cloud management plane that transforms your databases into DBaaS. It does so by extending this control plane to any environment where you run your databases; they can be on-premise, or in any cloud form – public, private, or hybrid – and on single – or multi-node clusters. No matter where your database is deployed, the control plane remains consistent. This greatly simplifies the mundane, but often labor-intensive tasks of managing the full lifecycle of the database, from deployment to management, patching, and replication. You can call an API and get database server with connection string with best practices from the Era catalog. As part of this control plane, your team can consistently enforce the right level of security over your data, ensuring that only the right people can view actual or masked data, and whether they are allowed to read or write.

On the horizon, Nutanix will extend Era to support multi-cluster and multi-cloud scenarios. Enterprises will gain the advantages of DBaaS without the restrictions. Nutanix Era is being designed to support all major enterprise databases providing the one-click automation that can run anywhere – in their on-premises datacenter or their public cloud of choice.

Author

Tony Baer, Principal, dbInsight

tony@dbinsight.io

Twitter @TonyBaer

About dbInsight

dbInsight LLC provides an independent view on the database and analytics technology ecosystem. dbInsight publishes independent research, and from our research, distills insights to help data and analytics technology providers understand their competitive positioning and sharpen their message.

Tony Baer, the founder and principal of dbInsight, is a recognized industry expert on data-driven transformation. *Analytics Insight* named him one of the 2019 Top 100 Artificial Intelligence and Big Data Influencers. His combined expertise in both legacy database technologies and emerging cloud and analytics technologies shapes how technology providers go to market in an industry undergoing significant transformation. His regular ZDnet *"Big on Data"* posts are read 25,000 – 30,000 times monthly.