Technical White Paper

4 Steps to Building an Enterprise Cloud

Tintri

@tintri www.tintri.com

Contents

What is “Enterprise Cloud?”

Given the multitude of opinions of what constitutes a complete definition of “enterprise cloud,” any serious discussion should begin with a level set on what that term represents. While definitions of enterprise cloud vary significantly depending on the bias of the presenter, NIST has provided a standard for the elements that make up a cloud computing technology.

On-demand self service: Automatic unilateral provisioning

Broad network access: Connections by heterogeneous clients

Resource pooling: Multi-tenancy resource pooling by provider

Rapid elasticity: Dynamic scale and re-provisioning

Measured service: Metered billing based on granular service units consumed

Pulling these all together into an integrated and coordinated infrastructure offering constitutes an enterprise cloud framework. But implementation paths may vary, depending on priorities and needs. So organizations may deploy these elements at different times and in a different order, in accordance with the demands of their users.

The user experiences all of this enterprise cloud infrastructure at high levels of abstraction. Geography, physical servers, network pathways, storage types, etc., become transparent and even meaningless. Conventional infrastructure concepts such as LUNs are replaced by virtual machines, containers, and vDisks. This means enterprise cloud architects must transform how they think about infrastructure management and adopt new perspectives on optimizing performance, cost, and risk. This is especially true for storage, since storage is used by and affects every other element of the cloud infrastructure.

When planning an enterprise cloud, there are fundamental considerations that must be weighed to ensure the result can meet all current and future needs. The following sections outline a useful thought process to guide the design of an enterprise cloud and cover key implementation options. As you progress through this process, Tintri storage streamlines your implementation with our flexible and robust enterprise cloud platform.

Step 1: Assess Application Needs

Infrastructure is built to meet the needs of the applications that are deployed to run on it. While public cloud is well suited to cloud native applications (e.g. DevOps and eCommerce), that may not be the case for legacy enterprise applications (e.g. servers and databases). Whereas a central benefit of enterprise cloud is its ability to support BOTH legacy enterprise applications AND cloud native applications.

So, in considering whether certain applications are candidates for public or enterprise cloud, it’s important to account for these characteristics:

- **Age.** If an application is more than 10 years old, its language may not be supported by public cloud. Legacy applications tend to have complex interface requirements and a long list of enhancements and support requests, making public cloud migration problematic—they will perform more predictably in enterprise cloud. Whereas cloud native applications can thrive in either public or enterprise cloud.
- **Data coupling.** If an application needs its data to be “local,” it may not translate well to public cloud, and should be considered a candidate for enterprise cloud.

- **Compliance.** There may be instances where an application is subject to regulation or compliance requirements that may not lend themselves to public cloud instances.
- **Team size.** For applications that are used only by an individual or small group in a single location, it may not make sense to migrate it to public cloud. For larger or more dispersed teams, public cloud can add significant value.
- **Complexity.** If an application is tightly integrated with other applications or requires specialized infrastructure, it is less suited to public cloud. These types of applications are often resource intensive, which can run up public cloud bills, and require a degree of resiliency that can only be offered in enterprise cloud.
- **The cost of "undo."** If an application is moved to the public cloud and the result is unacceptable, how easily can that application be moved back as a local instance? If failure cost is "high," then enterprise cloud reduces risk and cost.
- **Other potential factors:** Agility, sensitivity to exposing an API to the outside world, reliability and performance.

In the end, the primary drivers to migrate applications to public cloud are usually scalability/elasticity and agility/time to market. But rather than react to those requirements alone, it's important to consider the above characteristics—it may be that those applications can access the same agility as public cloud in an enterprise cloud, AND receive greater coupling, compliance and recoverability.

That's why as we consider the next steps we will specifically focus on enterprise cloud—delivering public cloud capabilities in your data center.

Step 2: Establish the Right Architecture

Once the decision has been made to move to the cloud, migration of certain infrastructures can be either a breeze or a migraine, both for setup and continuing operation, depending on both initial and ongoing decisions. It is critical that these decisions on configuration and management work with both the capabilities and the limitations of the cloud. Tintri Enterprise Cloud is ideally suited to provide storage resources ideally matched for requirements of the enterprise cloud.

Identify and Allocate Storage Resources: Storage is fundamental to the success of any infrastructure, so taking the time to identify requirements and then operationalize them is a critical first step to cloud set up. Tintri dramatically simplifies storage management in cloud environments.

Unlike localized applications with dedicated physical resources, cloud applications and their resources are virtualized and shared. When conventional storage or standard flash becomes part of a private cloud deployment handling virtual workloads, it brings tremendous management baggage. Conventional shared storage systems manage objects such as LUNs, volumes or tiers, which have no intrinsic meaning for VMs. Each new VM instance must be assigned a specific storage LUN or volume. When IO requirements and VM behavior are poorly understood, a painful trial-and-error process ensues to make sure the storage needs of each VM are met.

Performance-sensitive applications can end up competing with other applications for resources, resulting in some of those applications getting starved by noisy neighbors. Without resources that have intelligence to manage this problem, the only two options are to have applications slow and/or time out or to heavily overprovision resources to deal with any load. The former is inherently untenable. The latter can be prohibitively expensive.

Tintri avoids these noisy neighbor problems by eliminating LUNs and assigning every VM its own lane. Tintri acts as a single, federated (loosely coupled) pool. This greatly simplifies and provides broad automation of resource allocation in your private cloud environment. With Tintri, there's no conflict over resources or policies and therefore, no noisy neighbors.

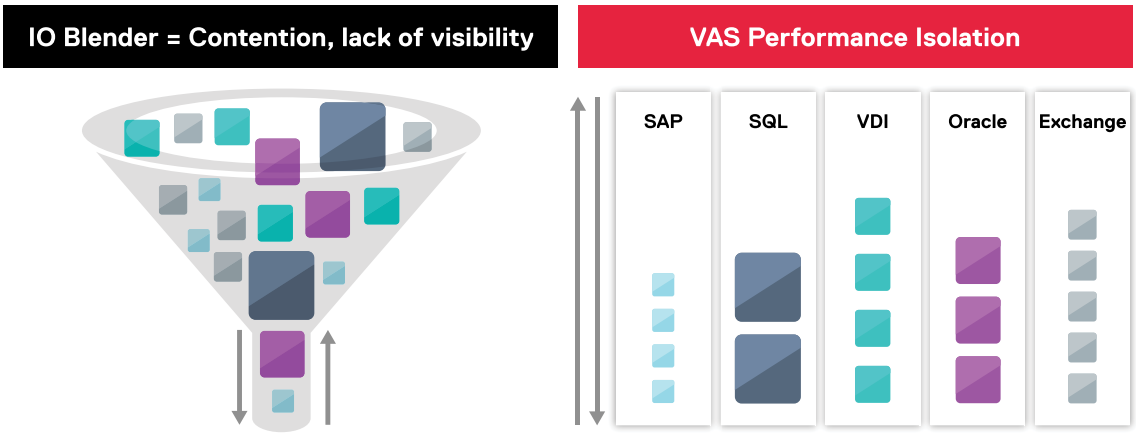

Policy Definition: After storage resources have been identified and an allocation schema put in place, it is important to establish policies for data protection/replication, QoS, and other operations and assign to service groups. How will IOPS be allocated per VM? How do you manage mixed workloads with different requirements? What process will accommodate different replication requirements of individual VMs?

With Tintri's advanced QoS capabilities, administrators can now allocate exact maximum and minimum IOPS to each individual VM, not just to a LUN or volume. Unlike conventional QoS, which requires administrators to predict the right IOPS values, Tintri provides visual guidance on the QoS values to specify, removing guesswork.

Tintri's granular QoS policies help you manage mixed workloads for different customers, with different service level requirements, while hosting multiple types of hypervisors on the same Tintri VMstore in a virtualized data center or private cloud. Tintri delivers true mutli-tenant virtual environments with support for chargeback.

Building a cloud around Tintri streamlines resource allocation and reduces the delays and risk of complex policy development and management.

Step 3: Automate

Once you have your storage resources identified, a strategy in place to allocate them among apps as needed, and developed policies for optimizing how they function, the next step is to think about how those storage resources will be accessed. To achieve the true benefits of the cloud, storage overhead must be minimized without sacrificing reliability and performance. Creating complex, event-specific scripting to pull and push data between apps and storage cannot scale and is prone to error. The way to avoid this is to build your storage and apps around a simple, comprehensive yet clean storage API.

This API should have granular command capability but this should also be coupled with the ability to bundle commands around a complete storage action. In other words, if meeting application data needs requires a series of individual puts/call API commands, the network chatter and administrative oversight can bring operations to a near halt. It is critically important to avoid repetitive, manual efforts by leveraging automation scripts and making entire storage operations one or just a few API calls to execute the required actions.

Ideally, APIs can be organized according to functional buckets for easy management. Creating API categories around such operations as provisioning and policy management greatly simplifies storage administration, even in complex cloud deployments.

APIs that are kludged together with redundant actions, inefficient organization and operation, and potential conflicts can, at best, reduce performance and complicate administration and, at worst, lead to time outs and down time. So, APIs need to be cleanly and efficiently constructed to allow infrastructure to be assembled like building blocks that easily snap together and share information.

Step 4: Analyze and Scale

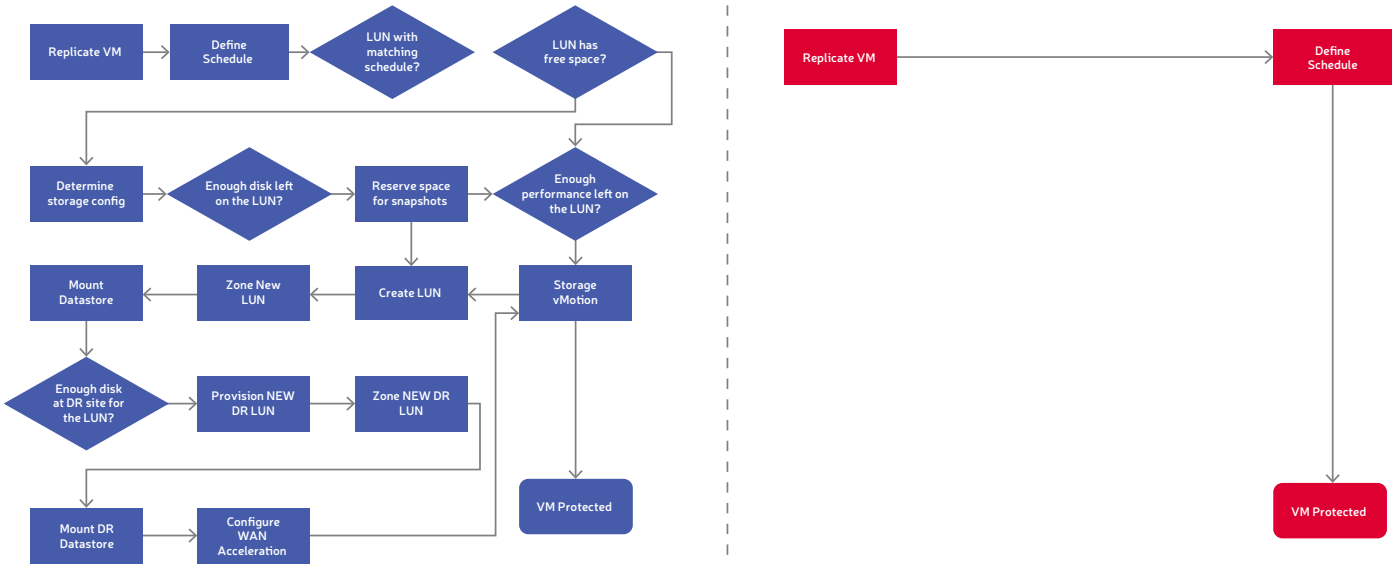

Ok, you've got your storage needs defined, policies identified, and orchestration via API outlined. The next consideration is ensuring you have a strong foundation for understanding and adjusting your storage as needed. This requires powerful yet easy-to-interpret analytics and a robust platform for scale.

Conventional storage's resource pool structure has no intrinsic meaning for VMs, making granular real-time monitoring for visibility on a per-VM basis and visibility into enterprise cloud latency issues all but impossible without massive additional effort. Granular visibility drives real-time and predictive analytics across cloud-based hosts, networks, and storage, addressing a significant shortcoming of conventional storage. For maximum value, these analytics must provide real-time actionable insights at a VM-level to rapidly identify factors that could contribute to increased cloud storage latency, degraded performance or even downtime.

Tintri's analytics provides customers with cloud-based real time and predictable analytics to improve data center planning and operations. These analytics can crunch millions of data points from 160,000 VMs over several years in less than one second. Today, Tintri can model virtualized applications' need for storage capacity and performance; an upcoming release will extend analytics to forecast need for compute resources in support of organic growth and simulated projects. Tintri:

- Optimizes your VMs based on a complete picture of their storage capacity and performance needs
- Gives you least-cost recommendations, saving you time, bandwidth and capacity, to maintain optimal VM distribution
- Allows you to review and edit the recommendations, and see the outcome before committing
- Learns every time you edit its recommendations and allows you to opt certain VMs out of migration

Analysis is only one part of the administrative challenge around storage for enterprise cloud applications. One of the big benefits of cloud is elasticity. But trying to manage conventional scale-out storage for virtual environments can be exceedingly complex. Here are some of the extensive storage management challenges with conventional scale-up storage.

- LUN- and volume-level data leads to bad guesses and poor decisions about optimal VM placement
- Poor migration recommendations based on incomplete data, and no visibility in to the impact on performance or the time required to complete a migration
- Problem VMs constantly getting bounced back and forth between arrays reacting to, but never resolving, performance issues
- Having to throw out your old storage, migrate to a scale-out architecture, and hire a team of Storage PhDs to run it

Tintri Enterprise Cloud is built around delivering a flexible yet robust foundation for your enterprise cloud. It directly supports all five elements of the NIST standard.

1) **On-demand self service.** Tintri VAS automatically abstracts away the complexities of cloud storage by providing a simple and intuitive interface that lets users take direct charge of meeting their storage needs (subject to administrative controls).
2) **Broad network access.** Tintri supports a broad range of clients needing access to enterprise cloud capabilities.
3) **Resource pooling.** Tintri presents and manages enterprise storage as a single, federated pool of capacity, enabling more flexible access and higher levels of utilization.
4) **Rapid elasticity.** Tintri scale out and automated provisioning ensures placement of virtual machines are optimized across the entire storage footprint.
5) **Measured service.** Tintri provides predictive analytics that show requirement trends and delivers forecasts to ensure needs continue to be met. With Tintri, organizations gain visibility across their infrastructure to help pinpoint issues both within their storage infrastructure and also those around networking and compute resources.

Tintri provides a simple yet powerful platform that serves as the foundation of your path to the enterprise cloud. With Tintri, transitioning to the cloud is both accelerated and flexible, allowing you to take whatever path makes the most sense for your organization.