# Cloud Ops: Operating OpenStack™ Clouds

## An Operational Approach to Deploying & Sustaining Anyscale Environments

A Dell technical white paper

By Rob Hirschfeld & Greg Althaus

**Revolutionary Cloud and Big Data Solutions**

The power to do more

# Cloud Ops: Operating OpenStack Clouds

**Bring open APIs and best practices to data center operations.**
**Tags: #CloudOps, #DevOps, #Hyperscale, #OpenStack, #Crowbar**

## Table of Contents

## Executive Summary

Dell's Data Center Solutions team (DCS) has helped deploy clouds for some of the world's largest and longest-running hyperscale data centers. This white paper shares some of the critical path issues that we encountered while working with our customers, and the resulting operational learning we gained along the way. We will walk you through how we have applied our experience and expertise gained in implementing the principles of "DevOps" for cloud infrastructure.

From the very beginning, we realized that we needed to go beyond traditional IT infrastructure and management if our customers were going to be able to respond successfully to rapidly changing cloud environments. Our learning can be distilled down to two key cloud attributes:

***"The cloud is never at rest. The cloud is always ready, never finished."***

We ended up taking an operational approach to help automate our customers' system configurations and management; we integrated tools and processes from software development methodologies into IT operations (often referred to as "DevOps"). To take the "Cloud Operations" model even further, we integrated it with an open source platform model (OpenStack™) and developed some additional tools (Crowbar) to help manage deployments. Cloud Operations delivered the operational and process controls our customers needed, and OpenStack + Crowbar has provided an openly accessible, scalable, and globally-supported development platform.

By the end of this paper, you will be familiar with the operational challenges of hyperscale clouds that we encountered, and learn about how we overcame these obstacles. Our goal is to help you become operationally ready to design your own trial system that will serve as the foundation of your hyperscale cloud. At the end of this paper, we have also provided a list of resources and references to help you tackle the technical aspects of building your own cloud from scratch, featuring recommendations based on our hard-earned expertise.

Cloud computing requires a shift in thinking.

*"We cannot solve our problems with the same thinking we used when we created them."*

*—Albert Einstein*

## Our Journey to Cloud

The Dell DCS group, OpenStack's initial sponsor at Dell and creator of the PowerEdge C line of cloud optimized servers, has been the supplier to the biggest data centers on the planet; consequently, we have had unparalled opportunities to learn with our customers and our customer's customers. The core lesson is that both running and using a cloud at hyperscale create fundamentally different challenges than anything else that we had managed.

Our journey to cloud computing starts with understanding these challenges.

Anyscale or Hyperscale?

While DCS' legacy is serving the largest customers, our mission is to take those lessons and apply them to customers running at any scale. By design, the same frameworks and operations models we deploy for small proof of concepts can scale up to the largest deployments.

## Challenge #1:

### "*How do you manage a cloud if you can't touch it?*"

Cloud data centers (our customers) were racking so many servers, switches, and drives that they had to rethink their core assumptions. Public cloud consumers (our customers' customers) had a nearly identical problem set because they were consuming virtual fractions of servers across the Internet.

Both classes of users had the same challenge:

| Problem | Cloud Operator Challenge | Cloud User Challenge | Solution |
|---|---|---|---|
| Lots of servers | Problems too big for single server | Virtual servers are limited | Scale out: Split problems. |
| Unreliable servers | Redundancy too complex/ expensive at scale | Weak SLA from cloud providers | Scale out: Move redundancy to software |
| Limited access | Cannot dispatch fast enough | All servers remote via API | Scale out: Move redundancy to software |

Both cloud groups see their applications as living organisms where speed is the key to survival.

## Challenge #2:

### *"How do you manage a cloud that won't stay still?"*

Cloud computing requires a different mindset to infrastructure than the traditional IT thinking. Because cloud computing is in a constant and rapid state of change, it requires a shift in approach to respond effectively and efficiently.

| Problem | Cloud Operator Challenge | Cloud User Challenge | Solution |
|---|---|---|---|
| Code is constantly changing | Deployment time > code delivery rate | Pace of innovation is fast | Continuous deployment |
| Resources are constantly changing | At scale, no system is steady state | Match output with consumption demand | Automated deployment |
| Too many components to easily track | Complexity is expensive | Require dynamic integration | Centralized configuration management |

## New Assumptions for Cloud

These challenges forced cloud operators and users to change our core assumptions. We learned to:

Cloud computing is...

Don't expect to find a formal definition of a cloud here. There are lots of flavors of clouds, and they are constantly changing.

Our focus in this white paper is to help you build and run your cloud successfully, regardless of what it looks like.

- Change from a few big atomic servers to many small servers (scale out). Getting a bigger server may not always be feasible or the best option.
- Change from hardware redundancy to software redundancy. Hardware redundancy was unavailable, too complex and/or too expensive.
- Change from manual configuration to automated deployment. Even with an abundance of rocket scientists, even really smart people cannot keep up with the pace of change and complexity.

In changing our assumptions, we went beyond traditional IT infrastructure and management. We needed an operational approach that balanced hardware, software, infrastructure automation (Ops), and best practices (Expertise).


Since Dell has the cloud hardware and OpenStack provides the software, the missing pieces all about operations: automation and practice.

For the rest of the paper, we will talk about how adding Cloud Operations software and practices complete the picture.

## Finding Common Platform Standards & Cloud Infrastructure

Even with our experience in knowing how cloud operations and users solve their problems, it was not clear how DevOps could solve all of the problems that our customers faced.

The more clouds we built and deployed, the more obvious it became that we needed additional operational support to resolve deployment roadblocks. We determined that clouds require an operational focus above and beyond what we saw in the market. We saw that:

- Clouds demand significant operational and process controls.
- Operational decisions drive hardware and software decisions.

***What was missing: accessible, common platform standards to build clouds and a cloud infrastructure (Cloud Operations) to manage them.***

We resolved these mission-critical issues of infrastructure management and platform standardization with an approach that intersected two major technology waves: operational automation (DevOps), and open source cloud (OpenStack™). By combining these, we found a practical approach to scaling cloud deployments, managing elastic resources, and embracing constant change. We are now in the process of taking the services and products that we developed for our customers on a custom basis, and are building them into standard, fully tested, packaged, supported and marketed products, including the Dell OpenStack-Powered Cloud Solution.

### What animal is your cloud?

This is a funny question but not an idle one. Thinking about it will help you understand your approach to solving problems.

For me, it's a colony of ants because they collaborate locally to accomplish larger goals. That means I'm always looking for distributed solutions and eliminating centralized decision making.

# Cloud Ops: Operating OpenStack Clouds

Before we talk process, we need to set the technology stage. While cloud software is not our focus here, the introduction of OpenStack is significant to this white paper because it enables a new generation of open operations. An open source cloud creates the opportunity for the community to collaborate in cloud and operational best practices.

Like other technologies going through the early development and adoption part of their life cycles, the 2010 cloud market was highly fragmented. We had no standard "Cloud Operations Manual" or common platform to help our customers build, deploy and run their clouds—and there was no industry consensus on what these were:

- Amazon EC2 dominated, but the platform was not public domain.
- VMWare was building on their enterprise base while XenServer focused on clouds.
- Start-up cloud niche players appeared in the cloud market space, but the broader, integrated community-based projects were missing.
- Service providers could not agree on APIs.

By the summer of 2011, the landscape looks different. OpenStack offers a public domain IaaS platform for multiple hypervisors, with multiple entities contributing, and APIs implemented by multiple providers. By providing a cohesive, integrated open source cloud solution, OpenStack has disrupted the cloud market in a fundamental way.

## Operations Matters

A cloud is composed of two sides: a user (public)-facing side with APIs, and an internal (private) IT infrastructure side that delivers services to consumers through these APIs.

The first side, the public-facing API + ecosystem, is where consumers interact with the cloud. The cloud ecosystem exists beyond the APIs, creating the cloud market, which is infinitely elastic. This half represents the services that the cloud provider delivers to consumers.

Whatever is exposed through the API to the consumer is the cloud provider's contract, and IT operations has to fulfill this however possible. Ideally, the APIs provide the smallest usable "surface area" to limit the provider's exposure, which frees the cloud provider to focus on scale. The"Ops Black Box" resides beneath the public-facing half of the cloud. Consumers have no visibility into this area. To the API user, it does not matter if the inside is ugly, inefficient, or inelegant—the provider just makes services happen for them.

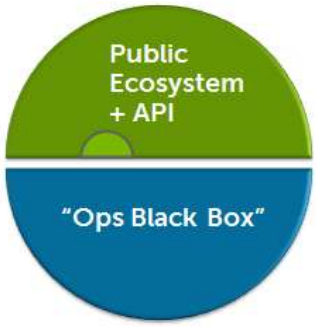


*Cloud =*
*Public Ecosystem & APIs + "Ops Black Box"*


*Consumers interact with the cloud through public ecosystem & APIs.*

***Except, of course, that it does matter a great deal to the cloud operator.***

### Cloud's Dr. Jekyll & Mr. Hyde

With apologies to Robert Louis Stevenson, clouds remind us of Dr. Jekyll and Mr. Hyde. Every cloud has two sides to it. With the public side, we expect it to behave mostly like Dr. Jekyll—helpful and generally well-behaved.

If the operations side ("Ops Black Box") is not managed properly, you run the risk of unleashing something closer to Mr. Hyde — unpredictable, dangerous, and a real pain to clean up after.

The purpose of this paper is to help you prevent the appearance of Mr. Hyde.

The operational half of the cloud should seem to be a perfectly performing, strictly partitioned and totally elastic resource pool. Historically, the "black box" portion of the pie has been further broken down into 3 key components, as shown in our next diagram ("Ops Black Box"):

- Hardware
- Software
- IT operations


*IT manages delivery of services in the private "Ops Black Box" half of cloud.*

As the diagram shows, the IT part of the cloud is more than just the hardware and the software. Without operations, there wouldn't actually be a cloud. We need processes and services. We acknowledge this as the reality of the situation. Historically, cloud providers have spent a lot of time talking about the left side of the pie, with very little discussion around right side of the pie. In reality, running a cloud is driven by operational processes and choices. The painful truth is if you want a cloud, you need a model for operating the cloud.

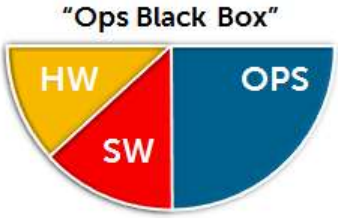


Cloud scale requires an operational foundation that can respond and react because deployed software and infrastructure is constantly evolving and adapting. Like the applications that run on the cloud, we need to recognize that cloud is a moving target and build operational systems that can handle that.

## What Is a "Cloud Operations Model?"

In the beginning, our first cloud offerings at Dell looked like this diagram—traditional IT installations. Literally.

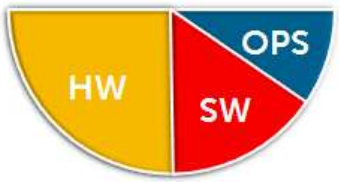


We would go to a customer to install their cloud. We would install the hardware components,set up the software required to run the services, and get the system ready for our customers to take over—basically set up their cloud and leave. And then our customers would start asking a lot of ongoing, operationally-related questions that had nothing to do with what we had just installed... What was missing was the additional and necessary operational infrastructure to keep the shiny new cloud running.

From these first installations, we learned a critical truth about the cloud:

***"The cloud is never at rest. The cloud is always ready, never finished."***

The hard, cold reality is that a cloud is a living thing, in a technology sense. It is a constantly moving, migrating, evolving, changing thing. What we were missing was the operations model that responds to "never at rest, always ready, never finished" cloud environments. This is really the operations model for a cloud. Yes, you can turn it on whenever you want; you only need to expose a tiny API. But you never are actually done.

We believe that it's not just what do you, but **how** you do it that counts.

The Operational Model is not a product. It is a process that is supported by tools and knowledge that is encoded into software.

If you're looking for more about process. We'd like to suggest looking up Eric Ries, Martin Fowler, and Mary Poppendiek.

All are excellent advocates for the types of processes that we have followed.

So we took a hard look at this statement ***"always ready, never finished."*** What would that mean? It meant that we needed a method of describing operations in a way that is consistent, that is repeatable, maintainable, with ongoing flow.

***To manage clouds, we applied cloud-proven DevOps tools to our physical infrastructure.***

What emerged was a completely new model. This new IT operational model utilizes the concepts of traditional software development, especially standards around change control and repeatable processes, and applies them to the operations system.

It goes far beyond the traditional IT role of being the system's technical owner/expert—these concepts become a part of the internal operations framework in order to respond to the constantly changing cloud environment. This is why and how what we and others are doing that makes this operational model different from more traditional IT infrastructures. (This is the origins the beginning of the cloud operations infrastructure model, or what Gartner Research and other industry players are calling "DevOps."[1])

## The Cloud Operations Model: A Different Approach to IT Ownership

Change drives the need for continuous integration and deployment. In Cloud Operations, principles of automation are applied to these processes. Operations and business processes become encoded as software. Different needs will govern different cloud approaches. This is not new, but the operational approach is done differently than other IT infrastructure models.

The Cloud Operations model approach is the opposite of the traditional release driven approach. It looks at the issue of change management, and says, "There too much change, too fast, so we will give our customers a flexible framework that creates repeatable, expandable processes to build their clouds and continuously move forward." Rather than attempting to manage control on behalf of the customer, the operations model creates a methodology that defines how the humans should interact with the system.

When you implement automation principles, you actually increase the proportion of the software piece even further compared to our original cloud deployment model, because the toolsets and changes that you use become embedded as software—the majority of it becoming operational software.

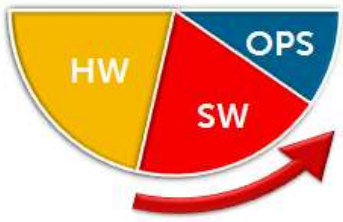


This additional software is not really product, but encoded business processes. In some ways, the Cloud Operations model ends up looking somewhat like a VIS model. The key difference from a VIS model has to do with the expectation and exposure of the product owner. In a VIS model, the product owner expects to use the product "as is," whereas the purpose of the Cloud Operations model software is to interact, modify, and build based upon changing customer needs.

We were not looking for an operational model.

We did not start out cloud building with an operational model.

We were struggling to deliver clouds the traditional way. It was solving the real-world problems that our customers brought to us that made us rethink our approach. Their challenges were not about installation, but about ongoing maintenance.

We believe that the operational model approach is the proven best practice for creating sustainable clouds.

*By 2015, DevOps will evolve from a niche strategy employed by large cloud providers to a mainstream technology employed by 20% of Global 2000 organizations.*[1]

*— Gartner Research*

[1] Cameron Haight |Gartner Research, *"Cloud Computing Will Change Next Generation Virtual Data Centers," slide 24 - "The Rise of a New IT Operations Support Model," 2010.*

## Key Tenets of the Cloud Operations Model

As a result of our cloud implementation experience, we have identified a number of key tenets to help manage large scale systems:

- **Change** —clouds are constantly changing, always ready. This requires a shift in IT strategy to be able to constantly and rapidly deliver change.
- **Continuous integration & deployment**—operations model utilizes constant system interaction, modifications, builds and deployments that are built around optimization needs.
- **Automation**—operations and business logic area encoded as software to minimize or eliminate the need for human touch. This creates a framework to create repeatable, expandable processes to build clouds and constantly move them forward. Ultimately, all the rules for operating the data center become encoded *as **automation software***.
- **IT Ownership**—IT proactively owns and interactively manages the cloud environment, using repeatable, expandable processes to build clouds and continuously move forward.
- **Layered deployment approach**—cloud systems require faster and more agile approach to deployments than the traditional "golden image" model.
- **Systems thinking**—whole system is viewed as an integrated unit, not just connected but isolated silos of hardware, software and operational infrastructure.
- **Cloud services**—operations model provides system-wide management functionality, but still delivers technology services as discreet components.

*"The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency.*

*The second is that automation applied to an inefficient operation will magnify the inefficiency."*

*— Bill Gates*

## Looking at the Cloud as a System

What makes operations such a big part of the cloud pie is that we are running a bunch of services. What we've been doing is automating the management of these services. In the past, in the traditional IT data center, one person did storage, another did networks, another did P&S, and another did VIS, with each person executing discrete tasks. What we are moving towards is a system-level thought on how all of these should be integrated together.

So we are seeing this more as an integrated system, rather than individual silos. Put these services into one operations model, but keep as discrete services. We want system-wide management, but the philosophy of services is that they are still discrete components. We don't want to blend everything all together, since in some cases that makes things even more difficult to manage. You want consumers and providers.

- **Systems thinking**—whole system is viewed as an integrated unit, not just but isolated silos of hardware, software and operational infrastructure.
- **Cloud services**—operations model provides system-wide management functionality, but still delivers technology services as discreet components.

## The Need for Speed: Layers vs. Images

Data Center departments have traditionally used a preconfigured image ("golden image") for deployments. This preconfigured image is a tested, known good configuration of all the components ***packaged together in a single unit***. This process is the current, established best practice for most IT groups. It uses the concepts of a library and cloning – you build a master image, deploy it, and then apply a configuration to it. Basically, you create a master, and then make copies of it.

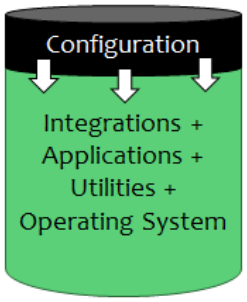


*Preconfigured master image used for deployments*

The main benefit to this image-based process is that you can push out large deployments easily, from one master to many copies. The downside of the image-based model is image sprawl.

To update or make changes using this model requires changing the whole image, and revalidating it. Changes have to be made to the master, and then branched out to all of the copies. In this scenario, the initial deployment is very convenient because of the large unit, but managing ongoing updates and changes becomes very difficult.

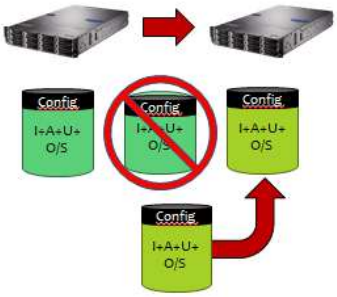


*"Golden" master image is pushed out one to many during deployment.*

The Cloud Operations model takes a different approach to deployment by using layers—there is no master, no "golden image," no cloning. The assumption behind this approach is that the cloud environment is ***"always ready, never finished."*** In other words, data centers are in a constant state of change, and IT must come up with a way to deploy and update in the middle of all of this.

Using the layered approach, we start with a minimal "image," and build up layers from this image. Controls are built in to manage layer functionality. It is harder to build the controls initially, but making changes and updates after the initial deployment is much easier, since you only update the layers impacted, not a whole large image.

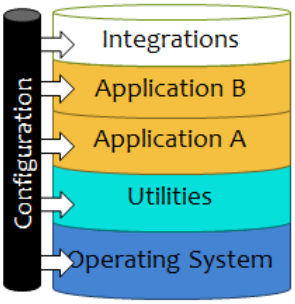


*Layered stack approach to configuration & deployment*

This tenet is what drives our Cloud Operations model. For IT organizations using the image-based model, it is very hard to deliver a unit of something and get it operational quickly. And because of this difficulty, there is usually a delay until the next round of updates and/or changes needs to be implemented.

This is a typical pattern for many companies—customers complained that the impact from these kind of frequent deployment changes was so disruptive that the change cycles for ongoing patches, release updates, etc. slowed down to a pattern of bigger releases done less frequently.

This model does not work so well in a cloud computing environment. Data centers are being driven to constantly deliver change. Pushing out single, massive units into deployment creates lag, and is difficult to manage. Cloud Operations takes the time to set up first—build the stack with the necessary controls, run integrations and other related

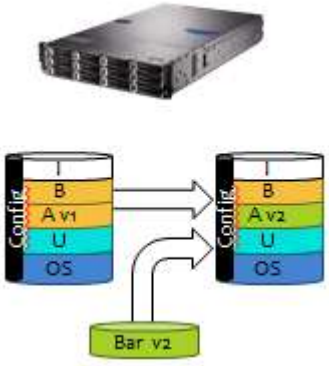


*Specific layers pushed out one to many during deployment.*

tasks. But once the layer is done, it can be deployed immediately, and that starts the continuous integration motor running. Once the minimal image set up is done, and a layer is ready to go, you can deploy. Changes can then be done on an on-going, as-needed basis—there is no need to hold back anything.

Implementing a layered model will seem really hard compared to the typical image-based model, and it is. If your assumption is that you don't deploy very often, this approach will probably the most painful. However, we are talking about an IT environment where things are changing so fast that there is a serious risk of not being able to keep up. There is no option to say, "I'm going to deploy THE image to my customers once, I'm going to do it every two years, and be done with it. I don't care about the new features." That does not work for the Web. The cloud is all about getting new features and fixes out on a timely manner, and it always has to roll out. It is all about the speed of execution, and continuous integration is crucial. The layered model is not only much faster than the image-based model, it is also more resilient. In the open source world of shared community development, this is an on-going, never-ending process, and it works well.

This is the reality of cloud space. Many of the largest Internet providers utilize the Cloud Operations model for their data center operations. For companies such as Facebook, Google and Amazon, deployment is a constant state of flux. These providers may actually have multiple versions of their applications running at the same time. They are continuously rolling out versions; there is no "don't change" lockout time. Sometimes they kill the version; sometimes they merge it into the main version. Their applications are constantly changing through very dynamic, organic, yet controlled/automated processes.

**Images vs. Layers**

| | |
|---|---|
| Like a compiled language | Like an interpreted language |
| • Early binding | • Late binding |
| • Tightly coupled | • Loosely coupled |
| • Assembled as a package | • Assembled in parts |
| **1-step** deployment | **Multi-step** deployment |
| • Big storage footprint | • Small footprint (packages) |
| • Cloned from "golden image" | • Setup per instance |
| • Single version | • Discrete versions |
| **Very static** operations | High change operations |
| • Manual setup | • Automated setup |
| • Low entry / high maintenance | • High entry / low maintenance |
| Works well in **enterprise** | Works well in **cloud** |
| • Hardware sensitive | • Hardware agnostic |
| • Brittle | • Elastic |
| • Typically: Windows | • Typically: Linux |

## Selecting a Cloud Platform

The cloud is not about selecting limiting technology, static end states, or single products. The cloud is about process innovation, business acceleration, and services. Having said that, different needs will govern different cloud approaches and platform choices.

We needed cloud Infrastructure that:

- Had support from major industry players.
- Was collaboratively developed without a single owner.
- Had an API that was service provider license friendly.

Additional reading

The focus of this white paper is on operations and software.

If you are looking for hardware-oriented information about provisioning hyperscale clouds, check out our white paper "Bootstrapping OpenStack Clouds" available at www.Dell.com/OpenStack

- Could be demonstrated to run at scale.
- Was built on open source components.

Dell has been at the forefront of hyperscale cloud implementations since the very beginning. What we have learned is that whether public or private, any hyperscale cloud environment requires a quantum leap in simplicity, homogeneity and density far beyond the capabilities of even the most fully virtualized enterprise data centers.

Based on our experience in implementing and provisioning some of the world's largest clouds, we have chosen to support open source platforms, specifically OpenStack, for a number of compelling reasons. OpenStack offers data centers:

- Massively scalable Infrastructure-as-a-service (IaaS) cloud solution that provides foundational components for cloud services
- Support for the top cloud application programming interfaces, also known as APIs (Amazon and Rackspace)
- Support for the top open source hypervisors (KVM and Xen)
- Options to run guests using Windows, Linux, or other x86-based operating systems
- Deployment at hyperscale (more than 1,000 servers) at multiple sites (NASA, Rackspace, and others)
- True open source—OpenStack has a significant, international base community that is constantly adding new features, allowing fixes, providing support, and extending features as needed

OpenStack is built on software developed by Rackspace and a broad, global ecosystem of partners. Dell is one of the original partners in the OpenStack community, which has now grown to more than 100 companies and over 2,000 developers around the world.

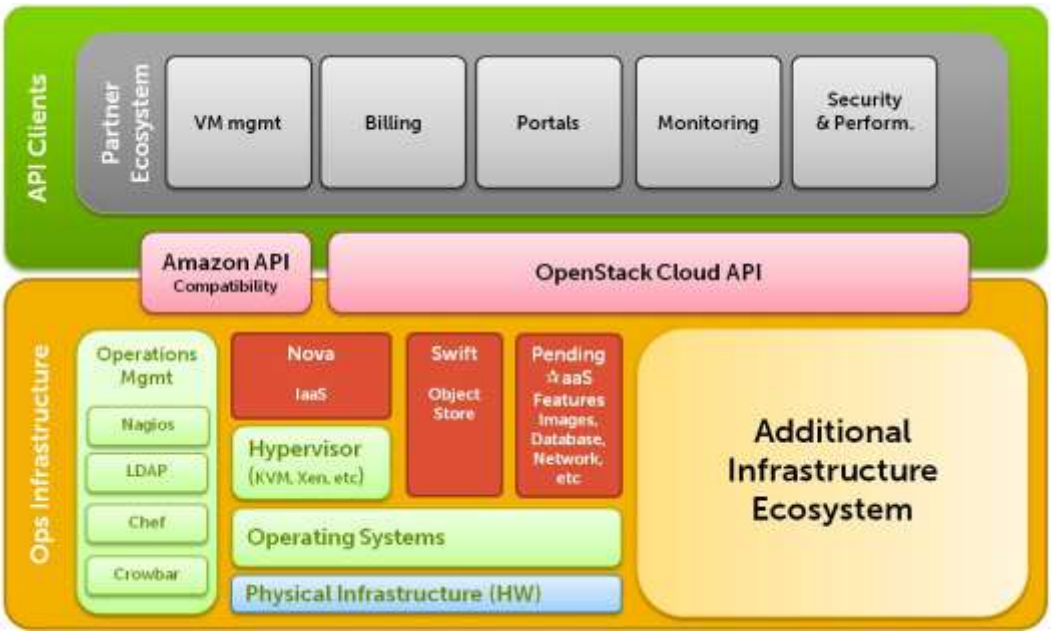

## The Idealized Cloud Data Center (Cloud Operations)

Many of our customers do not know how to operate a hyperscale data center because ***there is no commonly accepted pattern or operations model.*** That's where the Cloud Operations model comes into play—we have cloud-proven hardware and software, but no comprehensive description of the operational mojo.

With the evolution of the Cloud Operations model, data centers now have the tools and processes in place to respond to constant change with continuous innovation. The stage is set for the next phase: developing and implementing cloud computing best practices. From our cloud deployment experience, here are six of the key practices and characteristics of a successful cloud data center:

- **Assume change is ongoing and continuous**—Change operational models and processes to adapt to cloud environment demands.
- **Learn to deploy continuously and rapidly**—Adopt agile and iterative methodologies to facilitate frequent changes (avoid a "Big Bang" approach).
- **Utilize automation capabilities to minimize touch**—Undoing one mistake is painful enough, but having to undo the same mistake multiple times can be deadly.
- **Integrate operations management from the beginning, not as an afterthought**—Bake flexibility and agility into rules and business processes into the software ("infrastructure as code"). Integration is the key to success.
- **Focus on empowering innovation**—Successful cloud data centers learn from the mistakes of their predecessors, and are constantly innovating.
- **Active participation the Cloud Community/ cloud ecosystem**—Leverage use of custom and open-source code to reduce license costs and enable differentiation, get access to global cloud operations expertise through community engagement.

Back in 2001 when virtualization was a shiny new thing, no one had any idea on how to operate a virtualized data center. Companies struggled to win deals because our customers did not know how to operate virtualized hardware.

Ultimately, VMWare created the SAN-based data center consolidation pattern, and sales exploded. Companies adopted virtualization not only for the hardware (SANs) and/or software (ESX), but also because of the opportunity to standardize around an operations model they could understand.

## Considered Essential: Cloud Operations Tools

Cloud Operations tools should allow IT people to codify business rules and operations processes, making continuous deployments easier and faster. Using a collaborative software development approach, these new kinds of tools will allow IT organizations to apply development methods to infrastructure automation, saving time and resources through:

- Dynamic automation of core IT activities
- Repeatable, iterative configuration and deployment processes
- Massive scale of distribution to push change out

OpenStack provides an infrastructure foundation for hyperscale cloud; however, it is not a total solution. Depending on your objective, additional components will be required to operate your cloud. These components may enable software developers, integrate with internal systems, provide prebuilt templates, and extend customers' operations capabilities.

Some of components to consider are:

- Application support components include data storage services like structured databases (SQL), table storage (NoSQL), queuing services (AMQP), content delivery networks (CDN), and even an application programming platform (PaaS).
- Integrations such as billing, authentication, and VPN tunneling all help customers connect with their internal systems.
- Prebuilt templates and uploading images using open virtualization format (OVF) or similar technologies improves interoperability and allows customers to reuse work from other clouds.
- Operations services that take over customers' operations challenges by offering load balancers, firewalls, security services, backups, access monitoring, or log collection can be a substantial benefit for customers while leveraging your economy of scale.

There are an overwhelming number of opportunities to expand beyond the foundation provided by OpenStack. By investing in an open cloud infrastructure platform, you expand the ecosystem of services and partners. Having a shared platform reduces duplicated effort and having a large ecosystem encourages innovation and investment to solve difficult problems.

Note: Looking to get started with OpenStack? Dell offers the Dell OpenStack-Powered Cloud Solution, which includes a validated reference architecture, OpenStack and Crowbar software, and a number of key services.

Email us at OpenStack@Dell.com if you are interested in learning more.

## Productizing the Operations Model: Crowbar

To accelerate global adoption of the powerful OpenStack platform, Dell has worked to develop an effortless, out-of-the-box OpenStack experience with an OpenStack installer ("Crowbar") that allows bare metal deployment of OpenStack clouds in a few hours (versus a manual installation period of several days). Once development of Crowbar is complete, it will be released as an open source application to the OpenStack community.

Crowbar drives the Cloud Operations model to deploy components at ***all layers*** of the cloud taxonomy:

- All the components required to implement an ***entire*** cloud infrastructure are included, even ecosystems partners.
- Pluggable components deploy cloud infrastructure.
- OpenStack community and customers will be able to add to and expand the Dell IP.
- Can integrate with Dell existing products.

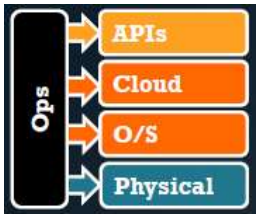


- Delivers basic data center services and required cloud infrastructure.
- Provisions bare-metal servers from box to cloud ***WITHOUT*** user intervention (other than racking/cabling and some minimal configuration questions).

## Key Takeaways

Designing a hyperscale data center requires thinking about operational problems differently. The large amount of resources creates complex management challenges, but also enables problem-solving by broadly distributing resources instead of relying on local redundancy.

Logical configuration is just as important as physical layout. Every step away from simplicity will cause exponential growth in complexity at scale. Find ways to automate and monitor.

To help accelerate evaluation of this powerful cloud platform, Dell has invested in ensuring an effortless out-of-box experience. Combined with Dell's tested and proven cloud-optimized hardware, our cloud installation automation helps ensure that you can confidently build a sophisticated infrastructure solution.

Dell is an active participant in the OpenStack community because we believe that OpenStack has the potential to bring our customers open APIs, capable practices to cloud operations, and affordable infrastructure.

## To Learn More

For more information on Dell OpenStack-Powered Cloud Solution, including whitepapers, case studies, and more, visit www.Dell.com/OpenStack or email us at OpenStack@Dell.com.

Visit:
Rob Hirschfeld's blog
The JB George Tech blog
Barton's blog / OpenStack