

# EMC IT ENTERPRISE HYBRID CLOUD

## Reference Architecture

### ABSTRACT

This white paper discusses the drivers and goals behind EMC IT's hybrid cloud transformation, the technical and business choices we made to enable the transformation, and the technology, process changes, and staff realignment involved.

June 2016



To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit [www.emc.com](http://www.emc.com), or explore and compare products in the [EMC Store](#).

Copyright © 2016 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see [EMC Corporation Trademarks](#) on [EMC.com](http://EMC.com).

Part Number H14757

# CONTENTS

|  |           |
|--|-----------|
| <b>INTRODUCTION</b> .....  | <b>4</b>  |
| <b>BUSINESS DRIVERS AND GOALS</b> .....                              | <b>4</b>  |
| <b>EMC IT HYBRID CLOUD STRATEGY</b> .....                            | <b>5</b>  |
| Defining cloud computing for the EMC IT solution .....               | 5         |
| Evaluating the existing data center environment.....                 | 6         |
| <b>PARTNERING WITH EMC ENTERPRISE HYBRID CLOUD ENGINEERING</b> ..... | <b>6</b>  |
| <b>EMC ENTERPRISE HYBRID CLOUD ARCHITECTURE AND COMPONENTS</b> ..... | <b>7</b>  |
| Core pod design and implementation.....                              | 8         |
| Automation pod design and implementation.....                        | 8         |
| NEI pod design and implementation .....                              | 8         |
| Tenant resource pod design and implementation .....                  | 9         |
| Disaster recovery for Enterprise Hybrid Cloud pods .....             | 9         |
| <b>ATLAS IMPLEMENTATION OF THE EMC ENTERPRISE HYBRID CLOUD</b> ..... | <b>9</b>  |
| Physical resources.....  | 10        |
| Software-defined data center .....                                   | 12        |
| Services library.....  | 13        |
| User portal.....   | 13        |
| <b>SECURITY</b> .....  | <b>15</b> |
| <b>CONCLUSION</b> .....  | <b>15</b> |
| <b>REFERENCES</b> .....  | <b>15</b> |

## INTRODUCTION

EMC IT's hybrid cloud transformation is designed to deliver a more efficient and agile IT data center environment and to offer Infrastructure as a Service (IaaS) to internal EMC IT customers—the business units at EMC. This white paper provides an overview of the EMC IT hybrid cloud project, internally named *Atlas*. It exposes the approach EMC IT took to transform a legacy IT environment into a self-service, automated, software-defined data center (SDDC) that can consume both private and public cloud resources. It discusses the drivers and goals behind the hybrid cloud transformation, the technical and business choices we made to create the hybrid cloud platform, and the technology, process changes, and staff realignment involved in the transformation.

With a small team of experienced architects, and leveraging existing data center resources, EMC IT designed and implemented a hybrid cloud platform that is based on the EMC Enterprise Hybrid Cloud 3.5 product. The architecture is anchored by the principles of cloud computing as defined by National Institute of Standards and Technology (NIST) (see [The NIST Definition of Cloud Computing](#)). The architecture also advances the philosophy of turnkey cloud computing that EMC and VMware are promoting.

Table 1 defines some of the terms used in this document.

**Table 1. Terminology**

| TERM                                | DEFINITION   |
|-------------------------------------|--|
| Atlas                               | The internal name for the EMC IT hybrid cloud project, which is based on EHC 3.5.  |
| Business group                      | In the context of Enterprise Hybrid Cloud, a group of users (often corresponding to a line of business, department, or other organizational unit) that access a shared pool of infrastructure resources and catalog items.   |
| EHC                                 | An abbreviation for Enterprise Hybrid Cloud that is used in diagrams in this white paper.  |
| Endpoint                            | A large pool of compute, storage, and network resources that is presented to the automation system. An endpoint can reside internally on VMware vSphere clusters or VMware vCloud Director instances, or externally on a VMware vCloud Air hybrid cloud.             |
| Hybrid cloud                        | A cloud platform that extends a private cloud into a public cloud facility. Network connectivity between the private and public cloud enables bursting or migration of workloads.  |
| Infrastructure as a service (IaaS)  | A standard set of automated resources that includes compute, storage, and networking capabilities that are provided through a hosting company or service provider (EMC IT in this case).   |
| Management pod                      | A generic term for the core pod or the automation pod, both of which provide core management services.   |
| Megablock                           | The EMC IT internal term for the consolidation of several small-footprint VCE™ VxBlock® systems into large endpoint resource pools.  |
| Pod                                 | An Enterprise Hybrid Cloud grouping of similar high-level services.  |
| Software-defined data center (SDDC) | A data center in which infrastructure is virtualized and delivered <i>as a service</i> . SDDC combines software-defined networking (for traffic management), software-defined compute (for workload management), and software-defined storage (for data management). |
| T-shirt sizing                      | An approach to virtual machine provisioning that offers virtual machines in predefined sizes (small, medium, large, extra large, and so on). Each size provides a specific, fixed number of CPUs and a specific, fixed number memory and storage.                    |

## BUSINESS DRIVERS AND GOALS

The legacy EMC data center compute infrastructure was designed primarily to support traditional second platform applications. These applications typically represent three-tier, scale-up designs, legacy storage allocation models, and Layer 2 network dependencies. This application model can be difficult to transform or migrate, and requires a high degree of manual intervention and expertise to manage and maintain. Second platform applications also assume a high level of resiliency in the underlying infrastructure.

Other challenges that are presented by the legacy infrastructure include:

- The infrastructure was not designed for automation or service abstraction.
- The way in which projects were instantiated, scheduled, and managed caused inconsistent object naming.
- Capacity management was not consistently governed as a single pool of resources.
- Few services were exposed through APIs.
- A variety of point solutions had been deployed, or developed, to accommodate deficiencies in product sets or rigid timelines. These solutions represented a high degree of customization.

Because of these challenges, we decided not to adopt a brownfield transformation approach. A system that is based on business policies, automation, and standard building blocks must be the rule and not the exception. The policies describe how data center components and services are expected to behave, and unambiguous policy definition is essential. The automation system must then be capable of applying the business rules using software algorithms. Policy automation is not easily accomplished. It requires a set of well-understood service objectives, as well as service attributes that can be accessed programmatically. The overall solution also must be stable, secure, fast, and agile.

## EMC IT HYBRID CLOUD STRATEGY

When our business customers first approached EMC IT to “build a hybrid cloud,” it was not immediately clear what such an undertaking would entail. The IT architecture team knew that cost savings and agility were top priorities for IT, but did not fully understand where, specifically, the existing data center design was deficient. IT had already successfully modernized the IT data center platforms with VCE VxBlock converged infrastructure (CI), and had completed conversion to a 100 percent virtualized infrastructure environment with VMware vSphere. These transformations had delivered significant operational and capital savings.

To address the requirements and concerns of the main stakeholders, the IT architecture team adopted a two-phase approach to defining and characterizing the cloud solution for the Atlas project:

1. Understand and define exactly what *cloud computing* meant in the context of EMC IT.
2. Evaluate the existing data center environment, determine what, if anything, must change, and then decide where we could achieve the most impact.

### Defining cloud computing for the EMC IT solution

The term *cloud computing* is one of the industry’s most popular, and often misused, buzzwords. To reach a common understanding of cloud computing for the Atlas project, we researched Gartner reports and respected websites and blogs. We also consulted with our peers in EMC and with outside consultants. This research and consultation revealed that some core principles are common to most definitions and implementations of cloud computing. These principles also align with the characteristics proposed in [The NIST Definition of Cloud Computing](#).

The NIST definition of cloud computing would be the pillar of our hybrid cloud transformation:

*“... a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”*

The NIST’s five essential characteristics of cloud computing would represent the building blocks for the IT services that we planned to provide to our customers:

- **On-demand self-service:** Self-service was already one of the capabilities most requested of IT and was foundational in our hybrid cloud design.
- **Resource pooling:** The ability to leverage a large pool of resources—be they compute, network, or storage—would be the next most important characteristic we embraced.
- **Rapid elasticity:** Aligned with resource pooling, and supporting a more agile response to consumer demand, would be the ability to add or subtract resources and components in an elastic manner.

- **Measured service:** Also of significant importance would be a means to affect consumption behavior by clearly exposing costs and usage metrics so that customers can make informed resource buying decisions.
- **Broad network access:** The final characteristic entails providing a seamless and ubiquitous network connection platform in which locality or device choice would not be significantly important.

To provide some guardrails and a beacon to follow during design and development of our hybrid cloud, we encouraged and evangelized these NIST characteristics in our proposals and designs.

## Evaluating the existing data center environment

With clear definitions for our new cloud system, we turned to the existing data centers to determine what was working and what must change. We concluded that resource isolation and high-touch custom installations were the two main issues that inhibited growth and stifled business innovation.

- **Resource isolation**

Although virtualization had resulted in considerable cost savings and efficiency, we found that we were slipping back to a policy of application and resource isolation, which was an approach that had been prominent in our pre-virtualization environments. Older run-rate budgeting policies and long-term project designs were partly the cause of this backsliding. The absence of a strong governance model and a full chargeback system was also challenging the ability to request and build virtualized environments quickly.

VMware environments were being rolled out as segregated clusters based on licensing, function, location, and business process. A common practice was to create clusters by environment (for example, development, test, and production), software (for example, Oracle and Linux), or tier (for example, mission-critical and support). The result was more than a hundred varieties of clusters, and cluster islands that are not easy to expand or share. The practice of segregating workloads based on prescriptive placement policies prevented best use of the total resource pool and precluded us from achieving the speed and agility that we require. If a business unit needed resources quickly, and predefined clusters were either unavailable or not easily expanded, IT had to limit or delay service delivery. This often inhibited innovation and our ability to react quickly to market demand. We had to address artificial resource isolation.

- **High-touch solutions**

An insistence on high-touch custom solutions was equally restrictive, and was often driven more by a lack of governance and a service catalog than by any business driver. With a static headcount and an increase in business service requests, we were unable to scale and support this high-touch model. It was evident that standardization and automation were critical to providing stable and secure services that could be rapidly deployed.

By understanding these deficiencies, and by planning a solution that would be based on the five NIST cloud platform characteristics and the EMC Enterprise Hybrid Cloud architecture, we were confident that our hybrid cloud approach would foster creativity, agility, and flexibility in support of a service-based IT organization.

## PARTNERING WITH EMC ENTERPRISE HYBRID CLOUD ENGINEERING

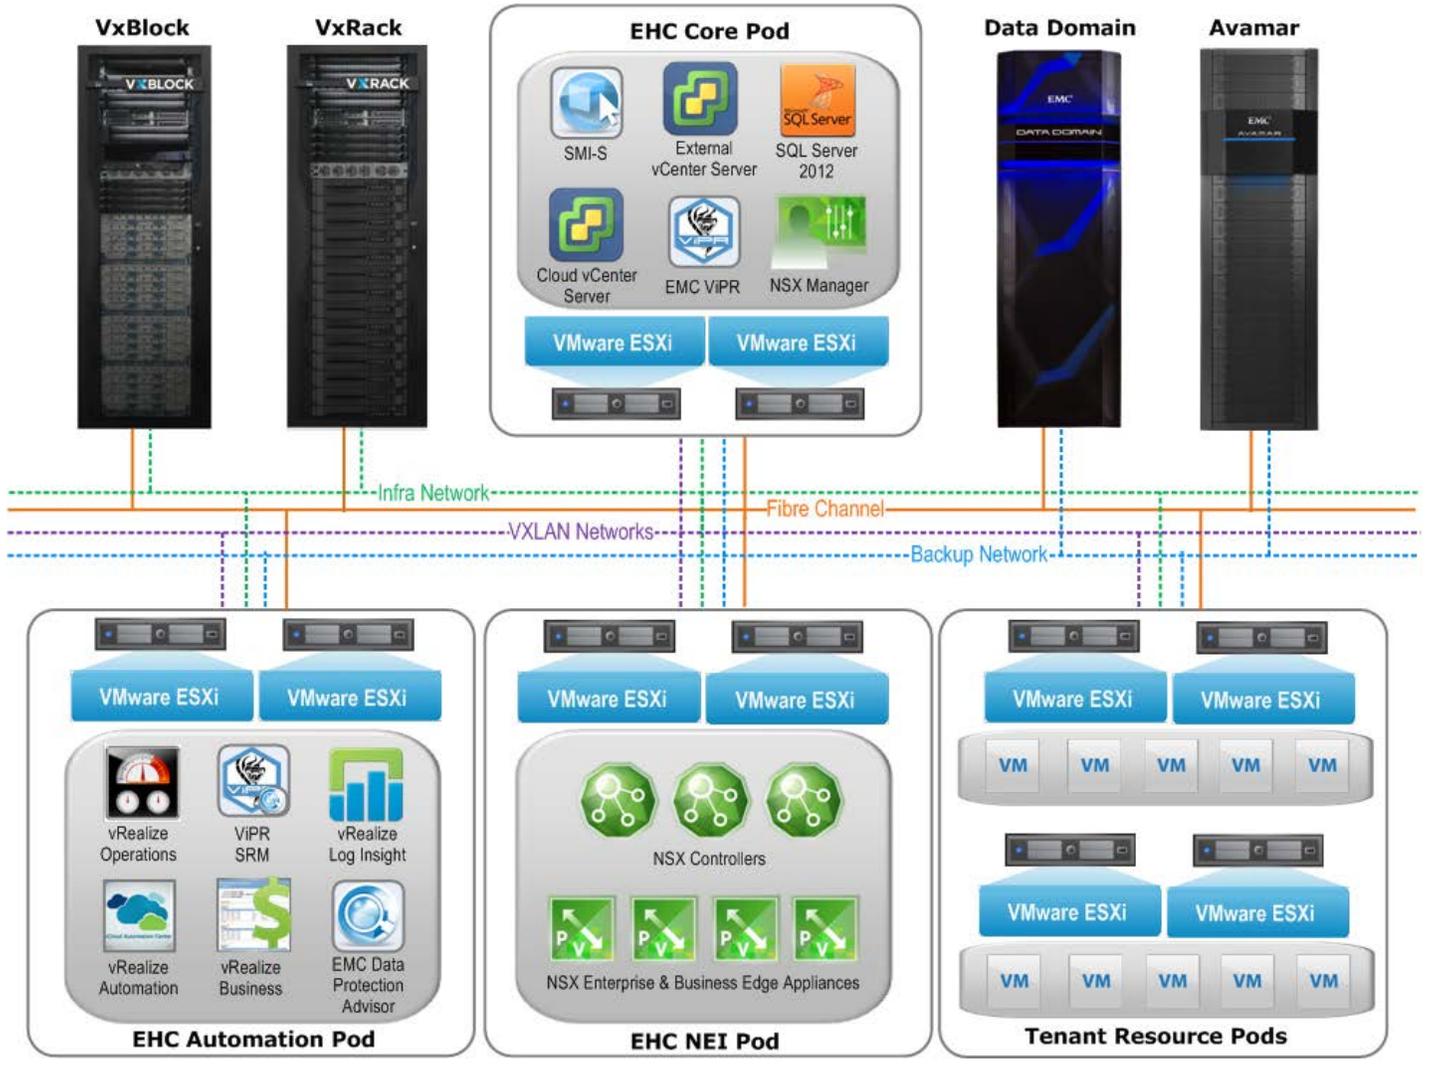
The EMC Enterprise Hybrid Cloud software stack and architecture design appeared to be perfectly aligned to where EMC IT wanted to go and to where other data center projects were already headed. However, we found several instances where our production environment did not specifically align with the Enterprise Hybrid Cloud support matrix. This lack of alignment was due to the commitment by EMC IT to use the latest releases of VMware products and because EMC engineering often calls on EMC IT to showcase the latest beta versions of EMC products in production. For example, our automation team had been in existence for some time and was using the latest VMware vRealize 6.0 products. We were also early adopters of the latest releases of the EMC VMAX3™ family and of EMC ViPR®. Networking was another area where IT had advanced the technology curve. VMware NSX 6.2 software-defined networking (SDN) was a key IT requirement from early in the project, although it was not fully supported in the initial release of EMC Enterprise Hybrid Cloud.

These early adopter initiatives provided an excellent opportunity to partner with Enterprise Hybrid Cloud engineering to ensure a rapid feedback loop and to help shape designs to meet enterprise customer expectations and requirements. We also used these initiatives to fast track EMC road map items and to provide our business customers with advanced cloud solutions.

# EMC ENTERPRISE HYBRID CLOUD ARCHITECTURE AND COMPONENTS

Figure 1 shows the architecture of EMC Enterprise Hybrid Cloud 3.5.

**Figure 1. Enterprise Hybrid Cloud architecture**



In this architecture, the cloud software stack is distributed across four main component environments referred to as pods (see Figure 1):

- Core pod
- Automation pod
- Network Edge Infrastructure (NEI) pod
- Tenant resource pods<sup>1</sup>

<sup>1</sup> Tenant resource pods are termed workload pods in EMC Enterprise Hybrid Cloud 3.1 and later.

Customers can deploy the core and automation pods on separate ESXi clusters, following the standard Enterprise Hybrid Cloud implementation, or they can combine the two pods on the same ESXi cluster. To enable the use of existing, excess compute power and to simplify the deployment, EMC IT decided to combine these two management pods for the Atlas implementation of Enterprise Hybrid Cloud.

## Core pod design and implementation

The core pod represents the core management domain. It hosts the base software components that are required to support any cloud deployment, including:

- The SMI-S and ViPR services, which manage and interface with storage array operations.
- VMware NSX Manager, which controls and manages our SDN service.
- The database servers for the cloud management stack.
- Two vCenter Server instances, which support services in the core and automation pods and the location-specific tenant resource pods respectively. For the Atlas project, we used vCenter Server Appliances to deploy both instances.
- VMware's authentication and authorization Single Sign-On (SSO) service, which is part of the VMware Platform Services Controller (PSC) component of vSphere 6.0.

Although we thoroughly understood the design and implementation of a standard single instance of the PSC, our requirements for a load balanced configuration at the front end of the SSO service presented some challenges, particularly with sticky sessions and certificate offload. The issues arose mainly because we were testing beta releases of NSX. To accelerate the release of these critical login services, we installed standard legacy hardware load balancing instead of the NSX software load balancer. With the recent release of EMC Enterprise Hybrid Cloud 3.5, a return to the NSX load balancer is scheduled.

## Automation pod design and implementation

The automation pod hosts all the automation components for the hybrid cloud solution. The pod's software stack represents the entire cloud management platform for the solution and the primary interfaces for administrators and users of the solution. This pod also hosts core EMC storage management services that expose APIs and provide software-defined storage (SDS) to cloud consumers.

The VMware vRealize Suite represents most of the software installed in the automation pod, including:

- VMware vRealize Automation and VMware vRealize Orchestrator, which provide the self-service portal and automation engine for all consumer service offerings.
- VMware vRealize Business, which provides cost modeling and chargeback, and which is natively integrated into the vRealize Automation portal.
- VMware vRealize Log Insight and VMware vRealize Operations, which deliver automated monitoring, performance management, and capacity optimization capabilities through APIs exposed by the traditional management and monitoring functions. vRealize Log Insight includes advanced logging, query services, and dashboards for analyzing unstructured data, while vRealize Operations provides similar capabilities for structured data and analytics.

## NEI pod design and implementation

The NEI pod hosts the core SDN routing and edge services of the NSX software stack. Theoretically, the NEI pod can be deployed on a shared management cluster. However, deploying the pod on its own ESXi cluster is the preferred option. A minimum of three servers is required in the cluster for a high availability (HA) configuration that can maintain a quorum. The provision of sufficient compute power, specifically memory, should be the focus of the NEI pod configuration.

**Note:** Network traffic that exits the cloud platform does so through the NSX edge services on the NEI pod.

For the initial deployment of Enterprise Hybrid Cloud, EMC IT placed the NEI pod on a set of blades in a large VxBlock CI domain. The proposed final location for the NEI pod is a set of rack-mount servers in an end-of-row configuration, with non-SAN storage

based on EMC ScaleIO. This location will enable IT to include the core networking service pod as part of a cold-start server environment with no reliance on other server management infrastructure or SAN support.

## Tenant resource pod design and implementation

Tenant resource pods represent the resources where consumer workloads reside and are processed. The Enterprise Hybrid Cloud design allows these pods to scale vertically within the limits of the computer technology stack and horizontally to include resources such as vSphere clusters, vCloud Air deployments, and vCloud Director environments.

For the Atlas project, we built the first tenant resource pod on a high capacity VxBlock system at our data center in Durham, North Carolina. These pods use the new tenant vCenter Server Appliance deployed on the core pod, as discussed in [Core pod design and implementation](#). We then built additional tenant resource pods at our data center in Hopkinton, Massachusetts, and at vCloud Air sites in Dallas, Texas, Sterling, Virginia, and Frankfurt, Germany. The resource pods at the Hopkinton data center have their own vCenter Server Appliance. vCloud Air resources do not include, or need, their own vCenter Server instances. Instead, they connect to vRealize Automation Center as endpoints.

Each of these pools of compute resources is an endpoint in the Enterprise Hybrid Cloud architecture. Each endpoint is managed as a separate pool that will have portions of its resources allocated to specific tenants of the cloud platform. The deployment of multiple endpoints enables us to easily add large amounts of private and hybrid resources to the cloud management platform, as required.

## Disaster recovery for Enterprise Hybrid Cloud pods

The core and NEI pods represent a site-specific solution set and do not need to be considered for inclusion in a disaster recovery (DR) configuration. Components with similar functionality are assumed to exist at the target DR data center. However, the automation pod is critical to the overall cloud management platform, which spans many data centers, and must be considered and addressed specifically in a DR protection scheme. EMC IT decided to use EMC RecoverPoint for Virtual Machines to selectively protect virtual machines in the combined management pod. Tenant resource pods that contain applications that require DR replication use storage replication and VMware Site Recovery Manager for this purpose. This particular DR service was not part of the initial Atlas release because the majority of IaaS requests that we receive are for non-production systems.

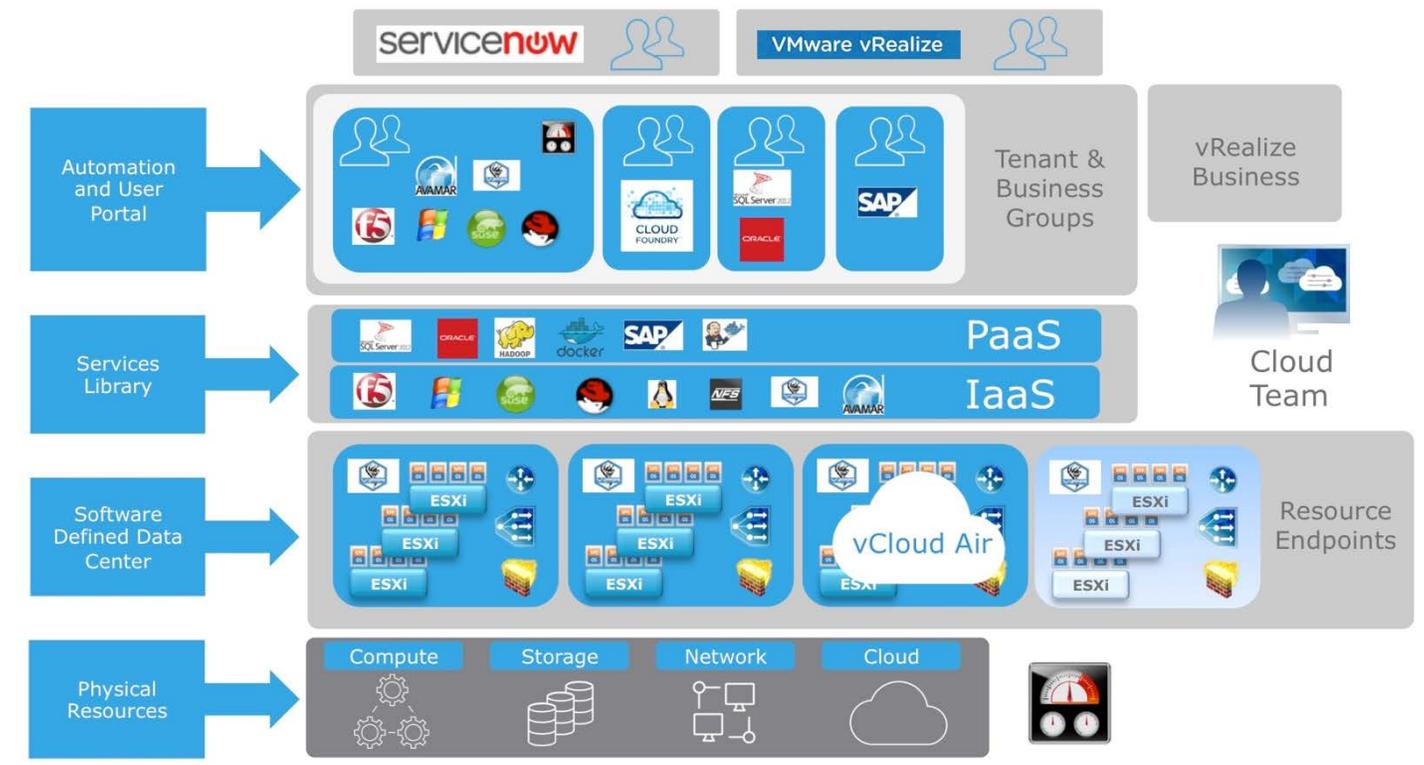
## ATLAS IMPLEMENTATION OF THE EMC ENTERPRISE HYBRID CLOUD

Installing and implementing the Enterprise Hybrid Cloud solution includes a highly prescriptive build process that is documented in a set of internal build guides. EMC professional services or a certified Enterprise Hybrid Cloud partner, not customers, normally carries out the process. Because EMC IT was an early adopter of the solution and a beta participant, we undertook the professional services role for the Atlas project. This circumstance, and because we were running beta releases of VMware software such as vSphere, NSX, and vRealize, introduced challenges and considerations that do not arise during a normal customer engagement. The base server infrastructure also presented a particular challenge because IT assembled the VxBlock systems in-house with early Cisco hardware and firmware releases. Consequently, both the Enterprise Hybrid Cloud and VCE teams had to validate our hardware and software environments before we implemented our hybrid cloud solution.

While the Enterprise Hybrid Cloud architecture provided the foundation for our hybrid cloud implementation, we developed the design further by leveraging cloud projects and processes that had already been launched within the IT organization. Figure 2 shows a high level view of the EMC IT Atlas implementation of the Enterprise Hybrid Cloud, which consists of the following logical layers:

- Physical resources
- Software-defined data center
- Services library
- User portal

**Figure 2. EMC IT Atlas implementation of the Enterprise Hybrid Cloud**



## Physical resources

The bottom layer of the technology stack consists of the core hardware infrastructure components—compute, storage, network, and cloud.

### Compute: VxBlock CI solution

Over the last several years, EMC IT has standardized on VxBlock CI solutions as the foundation for all data center server deployments. We continued this strategy for the Atlas project, though we are not restricted to CI solutions (see [Compute: Hyper-converged infrastructure solutions](#), [ScaleIO solutions](#), and [commodity servers](#)).

We anticipated significant demand for services after we released the self-service cloud platform. Consequently, the IT data center teams decided to convert and combine several existing, but small footprint, VxBlock systems into a larger consolidated platform, internally termed a *megablock*. These merged VxBlock systems include a switching fabric with a higher port count, over a dozen blade chassis, and a new VMAX3 array with service tiers. As customer demand increases, we can add any number of extensible VxBlock systems, each with up to 96 blades. We then use the blades to configure ESXi clusters based on vSphere configuration maximums, license requirements, or a demand for fault zone designs<sup>2</sup>.

At the start of the project, we evaluated the current trends in virtual machine sizing requests and mapped the trends against more than 11,000 existing virtual machines. We discovered that approximately 80 percent of all existing virtual machines were 2, 4, or 8 vCPU configurations. We also discovered that over 70 percent of existing virtual machine configurations were in the 2 GB to 16 GB range. Based on this knowledge, and on a forecast of more than 5,000 virtual machine requests over the remainder of the year, we established a standard EMC IT virtual machine compute size of 1 x 2.4 GHz vCPU and 2 GB of RAM.

We used this standard size to configure a new CI environment to support the more than 5,000 additional compute units that we expected to be requested. The sizing guidelines assumed no memory over-subscription for the virtual machines and a conservative 8:1 CPU over-subscription baseline. The guidelines also assumed a 20 percent additional performance gain from hyper-threading (HT) and a 20 percent hold back for high availability resource reservation.

<sup>2</sup> Fault zones are large, physically separated, virtual data centers, which are used by applications that require their services to be split across physically separate zones to provide high availability. The zones can exist within or across data centers.

The following is a summary of the sizing guidelines:

- 80 blades @ 20 core x 384 GB
- 1:1 memory overcommit
- 8:1 CPU overcommit
- 1,600 pCPU – 20% HA = 1,280 pCPU
- 20% HT = 1,536 pCPU
- 30,720 GB memory – 20% HA = 24,576 GB
- 8:1 CPU over-commitment = 12,288 vCPU
- Capacity (24,576/12,288 vCPU) = 12K 2 GB x 1 vCPU

Because sufficient resources already existed in our Durham data center, we decided to deploy the new VxBlock system in our secondary data center in Hopkinton to provide an additional tenant resource pool and automation pod. We proposed and ordered the specified VxBlock system with all the required base cabinetry, array frames, and switches, and with a starting set of servers and array drives. This configuration enables the capacity planning team to stagger capital expenditures by quarter to align with service request trends.

### **Compute: Hyper-converged infrastructure solutions, ScaleIO solutions, and commodity servers**

On-premises private cloud environments are not restricted to CI solutions such as VxBlock. We can easily incorporate hyper-converged infrastructure (HCI) solutions such as the VCE VxRack™ or VxRail™ systems, or any commodity rack or blade server configuration, into our hybrid cloud platform.

With increasing demand, and a shift to third platform and cloud native applications, we are finding that lower cost but highly scalable solutions, based on commodity servers and storage solutions like ScaleIO, are best suited for large scale environments. These platforms also fit seamlessly into our hybrid cloud design.

EMC IT has delegated the choice of server platforms to the infrastructure teams who manage vendor relations and to the capacity management teams who analyze usage trends and are responsible for data center component forecasting.

### **Storage**

Although the selected CI or HCI solution often predetermines storage solutions, EMC IT has traditionally leveraged VMAX arrays for pooled storage due to their advanced protection and service tier functionality. We continued this strategy for the Atlas project. For the first tenant resource pod in Durham, we selected a larger VMAX3 array to consolidate three older VMAX arrays. We also used a VMAX3 for the second tenant resource pod in Hopkinton.

Based on service levels and customer demand, we foresee increased demand for scale-out storage that weighs cost against enterprise service levels. For these workloads, we will deploy ScaleIO scale-out solutions that will be aggregated and presented to the automation and cloud management framework as just another resource pool.

### **Network**

The existing network hardware infrastructure that is used across the EMC enterprise data centers was relatively new and did not require any changes to support the cloud platform. The infrastructure includes core and access layer switches and firewalls, and delivers an architecture that aligns with the NSX SDN objectives of the Atlas project. Use of the NSX SDN software suite significantly impacted the network deployment by effectively lowering the operational and capital expenses related to expansion.

The network design enabled us to take advantage of Layer 2 encapsulation with VXLAN and to position workloads across data center regions without concern for routing complexities. From a connectivity perspective, all new CI or HCI solutions simply plug directly into the physical network switches. Edge services are available via SDN and provide ingress and egress (north/south) service between the virtualized hybrid cloud and the external or physical world. Distributed routers and the Border Gateway Protocol (BGP) are implemented to manage routing and east/west traffic.

The network design also enables extension of the Layer 2 networks to another data center if a large network outage occurs. However, only one of the two sites can be live, with default routing being advertised by the active site only. Because of this limitation, we use this configuration for DR environments only.

At the SDN layer, we enhanced data center security, and introduced further cost savings, by using NSX to implement network micro-segmentation. In the greenfield cloud platforms, all components are protected with a default deny-all policy that uses software-based firewalls operating at the hypervisor level. This configuration enables us to collapse private, semi-private, and DMZ functionality logically.

### **vCloud Air endpoints**

To enable IT to react quickly when data center expansion cannot meet timelines or cost profiles, we decided to extend the overall compute resource pool externally. To this end, we leveraged VMware vCloud Air to provide additional endpoints in Texas, Virginia, and Germany. These endpoints plug seamlessly into the vRealize Automation cloud management platform and enable easy horizontal scaling of the physical resource layer. The Texas site includes a Multiprotocol Label Switching (MPLS) connection that effectively extends the EMC data center networks into that facility. The site location provides a perfect DR target and a centrally located resource pool for west coast customers.

### **Software-defined data center**

The SDDC layer sits on top of the physical resources layer and is where the stability and efficiency of automation are most effective. Automation is facilitated by the APIs that components of the architecture expose. For example:

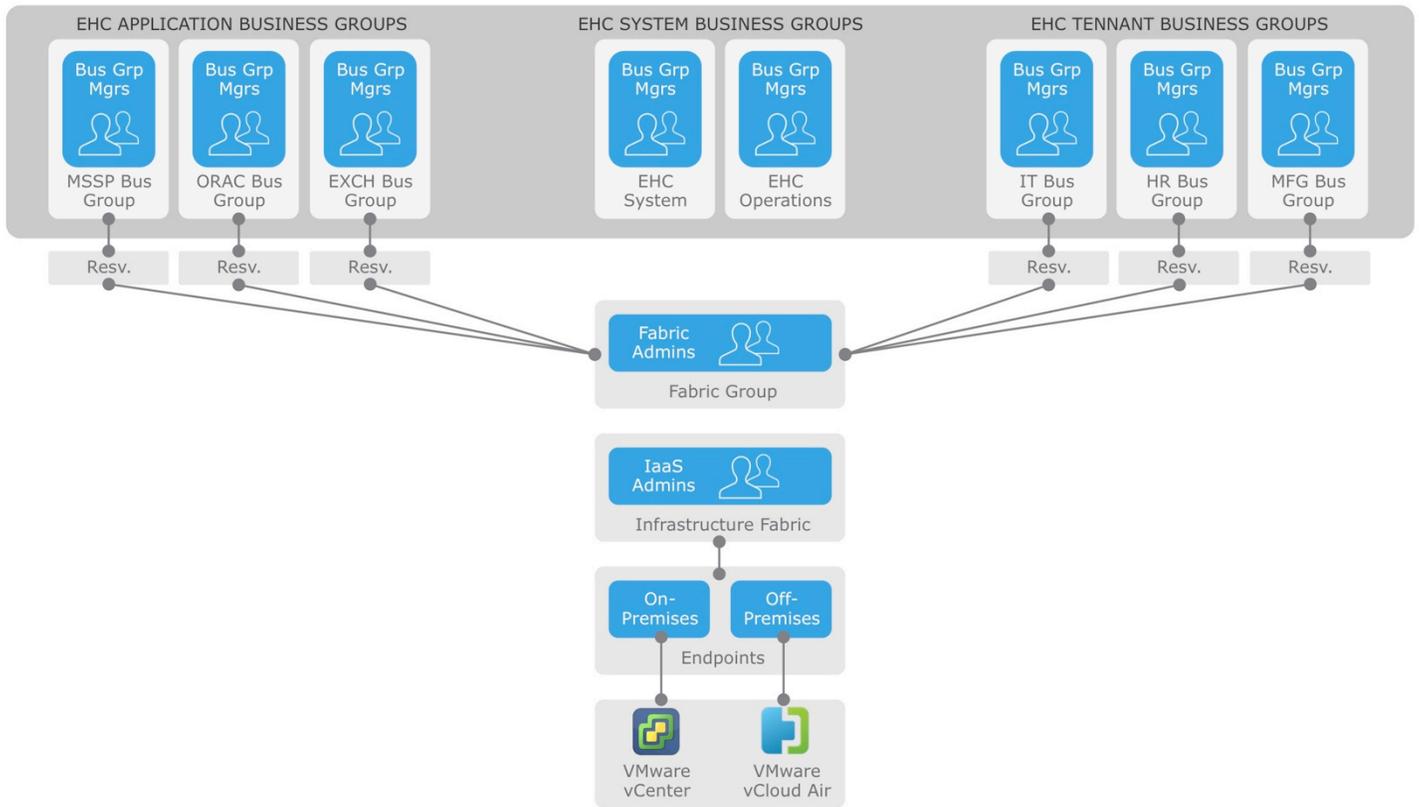
- All the standard components in the hardware layer include APIs that enable us to automatically assemble infrastructure components to deliver IaaS.
- vSphere APIs enable us to make programmatic calls into the SDDC layer from other software programs such as ServiceNow. With these APIs, we can modify cluster configurations and create, remove, or change virtual machines, networks, and storage resources.
- The NSX API exposes the capabilities of the SDN, enabling us to automate processes that normally require manual intervention. With this API, we can quickly and consistently manipulate and assemble services such as load balancers, routers, firewall rules, and switches through automation.
- ViPR abstracts the storage resources in the infrastructure layer and provides APIs that enable the automation software to create objects such as vSphere datastores and datastore clusters.

Because the hybrid cloud platform uses software to create compute, storage, and network entities, we can assemble these entities into larger pools of resources that vRealize Automation exposes as individual endpoints. vRealize Automation also represents vCloud Air resources as endpoints.

For storage administrators, the SDDC layer delivers a significant boost to provisioning speed, reducing the time that is required to allocate new disks to vSphere clusters from days to minutes. IT has provisioned over 1 PB of automated storage to date. For system administrators, the use of pre-defined code in the services library significantly reduces configuration drift and errors by automatically performing tasks that were previously labor intensive and prone to operator error. Also, fabric administrators can manage the resource endpoints through the user portal to allocate and reserve resources for customers by business group or function.

Figure 3 shows an example high-level view of the vRealize Suite tenant and business group design used by EMC IT.

**Figure 3. Tenant design and endpoints**



## Services library

The services library, which is the domain of the automation team, is layered on top of the SDDC. Using vRealize Automation and vRealize Orchestrator tools, the team develops the software blueprints and workflow code to create resource components such as virtual machines, datastore clusters, and load balancers. These components are then exposed as ESXi clusters behind the cloud vCenter Server systems in the core pod or presented as vCloud Air or vCloud Director resources. In turn, the ESXi clusters are presented to the cloud management platform as tenant resource pod endpoints, which customers are then empowered to use.

The automation team also uses the IaaS building blocks to create and expose additional layered services such as backups, IP management, operating system installation, application configuration, firewall modification, and many others. By using multi-machine blueprints, the team can also deliver platform-as-a-service (PaaS) offerings such as web/application/database configurations. Currently, over two dozen services are available in the library. In addition, the user portal includes interfaces to core services, such as vSphere, NSX, and ViPR, and to Infoblox (IPAM), Puppet, and ServiceNow.

The automation team operates as a Scrum team and delivers solutions to customer requirements in short, two-week sprint release cycles. In this way, IT can deliver products to customers quickly while maintaining the ability to iteratively improve on base services.

## User portal

The SDDC and the services library are of no benefit unless we can easily present services to our business customers. The user portal—a core component of vRealize Automation—provides a web service interface where consumers can be grouped and empowered to use the self-service portal to request IT resources.

We designed the user portal with a single tenant that includes multiple business groups. We assemble users into the business groups and use Microsoft Active Directory for authentication and authorization. The business groups typically correspond to lines of business, departments, or other organizational units. Splitting the portal into multiple tenants would provide a further level of isolation. However, multiple tenants require multiple vCenter Server instances, NSX domains, ViPR tenants, and so on, and increase the cost and complexity of a deployment. A single tenant supports the software maintenance, security, and isolation that IT required.

Each business group is allocated a portion of resources from the endpoints, which can reside across internal and vCloud Air clouds. Workload placement depends on the resource reservations that are allocated to the business group and is controlled by metadata that is evaluated during blueprint execution. Cost is the first criterion evaluated (is on premises cheaper?), next the criticality or sensitivity of the data (is it intellectual property?), and finally the functionality of the application (is it latency sensitive?). Each business group user is then authorized to run specific functions by using services from the service library. Figure 4 shows an example of the services presented by the user portal to a particular user.

**Figure 4. Atlas service catalog and a service request example**

The screenshot displays the EMC Atlas service catalog interface. The top navigation bar includes 'Home', 'Catalog', 'Items', 'Requests', 'Inbox', 'Advanced Services', 'Administration', and 'Infrastructure'. The 'Service Catalog' section shows a list of services on the left and a grid of service cards on the right. An orange arrow points from the 'Red Hat Server' card in the grid to a 'New Request' form. The form includes a table with the following data:

| Blueprint        | Description                    | Machines | Daily Cost |
|------------------|--------------------------------|----------|------------|
| * Red Hat Server | Red Hat Linux Server Component | 1        | \$184.00   |

Below the table, there are input fields for 'Request Information':

- Machines: 1
- CPUs: 2 (Select 2-32)
- Memory (MB): 8192 (Select 8192-98304)
- Storage (GB): 48

The primary users of the user portal tend to be IT employees who represent business groups consisting of database administrators, middleware developers, managed hosting providers, and storage services. However, integration with EMC's ServiceNow service management suite enables us to offer IaaS and PaaS services to the wider EMC user base without exposing service configuration details more suited to advanced technology users. Because ServiceNow can request services via the vRealize Automation APIs, we can expose general services, such as t-shirt sized (small/medium/large) virtual machines, to users through the ServiceNow portal. We can also leverage ServiceNow service, business, and operations management functions so that EMC employees can order services such as phones, conferencing, and laptops from the ServiceNow portal.

## SECURITY

Throughout the hybrid cloud design process, EMC IT insisted that security be a primary concern and be natively built in instead of being appended afterwards. To ensure that proper practices and compliance policies were followed, all architects on the project reviewed and responded to the Global Security Organization's questionnaire, which is based on the Cloud Security Alliance (CSA) [Cloud Controls Matrix \(CCM\)](#).

Active Directory groups protect and control access to the user portal and its services. All the original administrators with elevated access to virtualization, firewall, and storage tools are now restricted to using only the services for which they are authorized and to accessing those services exclusively through the user portal. Also, access to all hybrid cloud vCenter servers is rigidly restricted to a handful of senior staff only, in a break-the-glass model.

## CONCLUSION

The process of designing and implementing a hybrid cloud platform is challenging and requires a team of dedicated and experienced IT architects. Deploying a pre-configured, converged infrastructure solution like VxBlock, with a pre-installed and tested cloud software stack like EMC Enterprise Hybrid Cloud, greatly simplifies and accelerates adoption of the hybrid cloud.

Go to [Enterprise Hybrid Cloud](#) to learn more about the EMC Enterprise Hybrid Cloud product and solutions.

## REFERENCES

- [EMC Enterprise Hybrid Cloud 3.5: Foundation Infrastructure Reference Architecture Guide](#)
- [The NIST Definition of Cloud Computing](#) (NIST Special Publication 800-145)
- [Cloud Security Alliance Cloud Controls Matrix \(CCM\)](#)