



Case Study: NASA JPL

Truly Mission Critical: NASA's JPL governs their self-service hybrid cloud with Scalr.

Scalr provides the JPL IT Directorate with multicloud elasticity, resiliency, and accountability for OpenStack and AWS.

NASA's Jet Propulsion Laboratory IT Directorate carries a challenging mandate. It sits at the triple point of non-negotiable system demands of space exploration, fans rabidly hungry for the latest Mars Opportunity updates, and fickle government budgets. To meet their mission, JPL chose Scalr to manage their most critical on-demand cloud infrastructure. It powers their model-based engineering platform, guidance and navigation systems, prosaic shared Web hosting, and other services.



About NASA Jet Propulsion Laboratory

Renowned as simply "JPL", the NASA Jet Propulsion Laboratory is a US federally funded R&D center for operating planetary spacecraft and for Earth-orbit and astronomy missions. It operates the notable Curiosity and Opportunity Mars rovers, the Cassini-Huygens Saturn mission, Dawn, the probe investigating the asteroid Vesta, and the upcoming Europa mission to orbit Jupiter's moon Europa. Igniting the imagination of millions of space fans around the US and the world, JPL relies on an IT infrastructure that satisfies an "Internet scale" audience and the missioncritical requirements of space exploration.



Figure 1. Scalr allows JPL to offer internal customers a catalog of compliant services.

IT Infrastructure Needs

The IT Directorate's operating model is similar to many central IT organizations in large companies: cost-based chargebacks plus overhead. Their ultimate measure of success is internal adoption of their programs and services. The entire lab structure is a collection of mutually-interdependent service offerings like Guidance and Navigation platforms, Model-based Engineering platforms, radiation simulations on high-performance compute nodes and prosaic web hosting for public access. The IT Directorate is responsible for the critical underlying IT services for these platforms. These include network, cyber

"We looked at nearly everything. Scalr gives us what we need for a very modular infrastructure with flexibility and no lock-in."

– Chief Engineer
at the IT Directorate, JPL

security, storage, compute, policy compliance, shared hosting and others. For cost and technical reasons, special emphasis is placed on flexibility and the widest possible range of capabilities while being efficient with taxpayer funds.

To meet these needs, in 2009 JPL was an early adopter of the ondemand flexibility of cloud providers, using two major suppliers. In summer 2014 it has added an internal private cloud based on OpenStack. Shortly after adopting the public cloud, JPL published its best practices for creating a high-availability, fault-tolerant cloud service, coining the phrase “Cloud Oriented Architecture”. Being cross-cloud and hybrid public/private was an explicit goal in order to maintain vendor neutrality and agility.

To offer a pragmatic interface to their users JPL identified several technical requirements. First it wanted to create a uniform framework of APIs and cloud concepts that mapped to the various vendors’. Second, JPL wanted unattended restart and autoscaling services based on monitored events. That requirement mandated an orchestration engine capable of working across cloud APIs in different geographies and datacenters. Third, a business model of chargebacks required cost allocation with as little manual tagging as possible. Finally, JPL wanted to give developers a simple command-line experience to provision, deploy, and scale servers.

JPL Hybrid Cloud needs

- Uniform, cross-cloud API
- Public / private hybrid cloud
- Unattended restart
- Autoscaling and bursting
- Event notification
- Account cost assignment
- Simple developer management tools

The JPL Scalr Solution

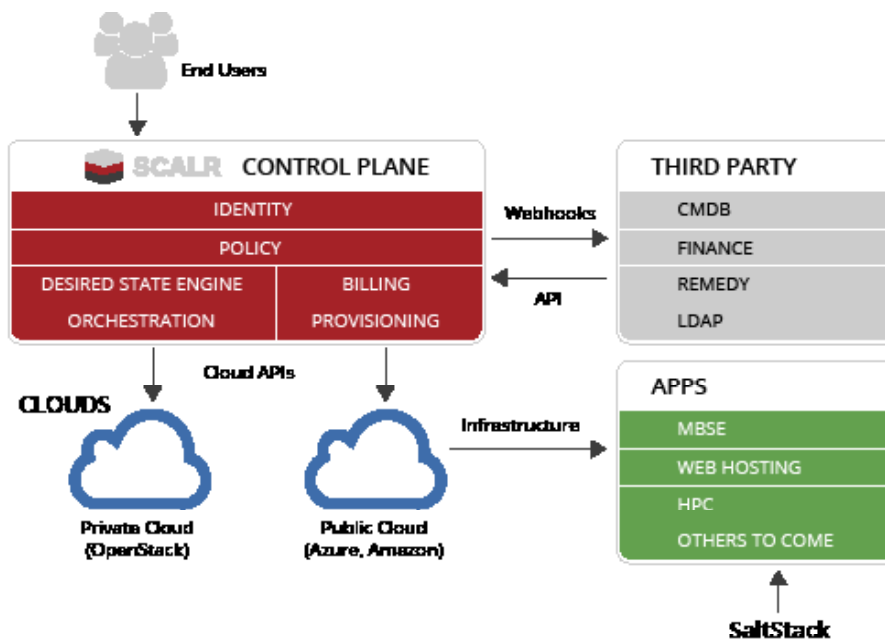


Figure 2: JPL's Cloud Architecture. A Scalr server instance is self-hosted to manage AWS and a private OpenStack cloud. It monitors events from managed resources, tagging them automatically for accounting, and the Scalr Orchestration Engine securely controls them. A common API maintains an abstraction barrier to simplify cross-cloud management.

After evaluating many hybrid, cross-platform cloud management systems, the IT Directorate easily settled on Scalr. It now provides the basis of their cloud platform. The JPL's IT Chief Engineer found that Scalr had all of the necessary features and was completely open and open sourced, leaving his team to feel confident for building atop it for the future.

Hybrid cloud

Scalr delivers a uniform interface across AWS and OpenStack which JPL uses for their main public and private cloud stacks. The uniform interface lets JPL enforce a common set of access controls, policies and management practices. This greatly reduces duplicate work and minimizes the chance of errors. Other cloud providers are also supported by Scalr like Google Compute Engine and CloudStack and Rackspace, providing options for future platforms. Azure support is coming Summer 2015.

Orchestration

Scalr's Orchestration Engine provides the core automation runtime for handling event-based, manual, or scheduled state changes to the cloud infrastructure. Along with SaltStack (using the reference architecture found at <https://github.com/scalrtutorials/scalrsaltstack>) JPL uses the Orchestration Engine for some of their most important functions: multi-cloud, unattended restarts, autoscaling, bursting, and management.

Environments and Interfaces

Environments are loosely isolated resource pools in Scalr's model. Cloud agnostic APIs provide an abstraction layer to the underlying resources like CPU, storage, and network.

Financial Responsibility

Scalr enables maximum financial responsibility for infrastructure at JPL. With multiple projects simultaneously demanding resources, the danger of runaway costs is extremely high. The solution is the showback of actionable cost information based on historical and projected cost data. JPL users are empowered to make financially sound infrastructure decisions thanks to the costs surfaced in the Scalr interface. Should JPL's needs expand to more powerful cost analytics and reporting, they can integrate Scalr with third party billing systems and rely on Scalr's Tagging Governance to supply the third party software with context.

Future Work

Given the successes with Scalr, the IT Directorate has its sights set on automating their HPC clusters using Scalr's templating and autoscale capability. Radiation simulation is one application that is "embarrassingly parallel". By using a Scalr farm with pre-compiled simulation software, the CPU load-aware events can trigger autoscaling up and down in the HPC farm. This shared dynamic system would replace a much more separated static allocation of HPC. The existing system is heavily underutilizing HPC resources or hampering researchers' experiments by taking too much wall-time. Exploration of the use of Docker containers is also underway.

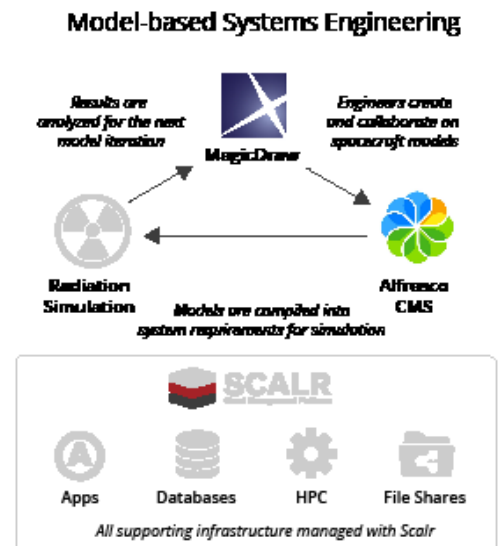


Figure 3

Example: Europa Model-based engineering

Figure 3: Model-based engineering is used to communicate between engineers in building the Europa mission. For example, Magic Draw and proprietary tools turn SysML into system requirements hosted in an Alfresco CMS. File shares are also available to design and engineering teams. These mission platform services are built atop the Scalr-managed hybrid cloud run by JPL's IT Directorate, creating servers, storage, and databases on demand.

Summary

Scalr gave the IT Directorate the capability to manage over 150 AWS and over 40 OpenStack developers' infrastructure with a tight, focused team. With hundreds of missions and platforms at JPL depending on this cloud, Scalr is part of a mission critical infrastructure for the US's space program. Scalr is honored to be a part of the US' most advanced research and infrastructure. Please contact us at info@scalr.com if you would like to learn more.

Solution Summary

NASA JPL Need	Detail	Scalr feature	Benefit
Infrastructure resiliency	Hybrid cloud unattended restart and relocation when Amazon reboots a hypervisor or other disruptive event.	Cloud event orchestration	Reduced operational overhead and panicked fire drills. Durable end customer and web browser experience.
Infrastructure elasticity	Hybrid cloud autoscale based on CPU utilization, or on one-touch manual demand.	Cloud scaling	Rapid response to bursty demand (e.g. Reddit post of Mars Opportunity imagery).
Public/private cloud hybridization	Ability to use public clouds for elasticity and private clouds for control with a single interface	Cross-cloud API	Enables choice between flexibility (public) and security/governance (private) or other reasons.
Multi-public cloud API	Ability to select more than one public cloud provider, with a single API	Cross-cloud API	Provides multi-vendor data center optionality for redundancy and pricing power.
Avoid Vendor Lock-in	Agility in purchasing and future-proof choices.	Fully Open Source.	Preserve future options. Maintain flexibility. Lower operational costs.
Downstream cost visibility	Chargeback internal customers want to predict infrastructure costs.	Cost analytics.	Provides visibility and predictability to internal end users on expected costs.
Multi-user cost tagging	Automatic master account to multi-user cost assignment.	Automatic resource tagging.	Applies automatic discipline to CPU, disk, etc. tagging. Avoids cloud cost sprawl.
Infrastructure modularity	Multi-tier and matrixed architecture including management and governance.	Cloud Orchestration and API. Fully Open Source.	Select best of breed in a highly dynamic marketplace. Ability to add/improve components with minimal disruption.
DevOps simplicity	Simplified, cross-cloud, command-line interface to developers, ala Heroku or ElasticBeanstalk	Cross-cloud API	Improve developer and ops productivity. Eliminate cross-cloud API training.

About Scalr

The Scalr Enterprise-Grade Cloud Management Platform enables today's enterprises to achieve cost-effective, automated and standardized application deployments across multi-cloud environments. Scalr utilizes a hierarchical, top-down approach to policy enforcement that empowers administrators to find the balance between the needs of finance, security, IT and development teams. Founded in 2007, leading global organizations have selected the Scalr platform, including Samsung, Expedia, NASA JPL, Gannett and Food & Drug Administration.