

Nimbix Federated:

A comprehensive architecture for secure, cloud-managed, multi-site supercomputing-as-a-service

Leo Reiter
CTO, HPC Cloud, Eviden

EVIDEN



Overview

High-performance computing (HPC) is a complex ecosystem that is increasingly being democratized by cloud and cloud-like technologies. This increased accessibility enables engineers and scientists to consume highly advanced resources without extensive IT expertise, advancing the digital transformation of the physical world. Now that industry has had a taste of on-demand supercomputing, the hunger for faster, more capable systems is growing. What has become clear is that public hyperscale clouds - while incredibly capable - simply cannot match the performance and raw throughput of the world's premier supercomputers.

For industry to continue to enjoy the benefits of supercomputing, it needs increased access to top-performing supercomputers, which comes with high direct and indirect costs. Besides large capital investments (or lease commitments), users looking to leverage HPC must recruit talent, build practices and processes, and suffer extraordinarily long lead-times due to global supply chain bottlenecks that simply can't keep up with the increasing demand for advanced computing components and systems. This translates to months - or even years - of delays that severely limit ROI as it is measured by private enterprises. That's assuming a successful roll-out, which is no guarantee.

Even the most expert supercomputer builders and operators suffer setbacks regularly! At its pinnacle, HPC pushes the envelope of physics so severely that failures are a fact of life, both at the deployment stage as well as the ongoing operational stages. Within a few years, it's time to do it all over again, because the next generation of supercomputers are so much more compelling. Riding this giant hamster wheel is a labor of love for those skilled in the art, but presents serious challenges and major barriers to entry for pragmatic industrial users looking to get ever more complex work done faster. Cloud computing drives a compelling bargain for many use cases, even in the high-performance space. Users have access to as much compute as they can deploy (or afford) at generally competitive on-demand rates. Unfortunately, when measured in terms of cost per job, diminishing returns on scalability mean that the numbers don't always make good business sense.

Additionally, while cloud storage is cheap and plentiful, high-end storage is not. Customers pay dearly for increased input/output operations per second (IOPS) and advanced software features like parallelism. Plus, the availability of high-end cloud storage, compute or networking is far more limited than the general-purpose tiers. This means that users often must wait for resources or settle for fewer than they need. Once again, the supercomputer is worth a good look, despite its complexity and cost.

Thankfully, governments and large organizations continue to invest heavily in HPC research, with an eye toward accessibility for the private sector. For example, the European High Performance Computing Joint Undertaking (EuroHPC JU) is focusing on building a world-class supercomputing ecosystem. Elsewhere, including in the United States, large direct investment continues to flow from massive funding bodies to advance scientific research in fields such as climate, healthcare and energy.



EVIDEN

BullSequana XH

EVIDEN

Qaptiva

What's unique about many of these initiatives is that they require participating centers to open a portion of their capacity to industrial users. In theory, this shouldn't pose much of a problem, because HPC operators provide users with access to systems day in and day out. In practice, however, there are major challenges, including:

- **Governance around data and geography**
- **Skill gaps between industrial users and research users**
- **Software licensing**
- **Service-level agreements (SLAs) that private businesses expect from providers**
- **Proper utilization, accounting and monitoring**
- **Billing**
- **Compliance with various rules and regulations for industry and government-funded research institutions**

There are countless other challenges as well, each of which affects end users and operators alike.

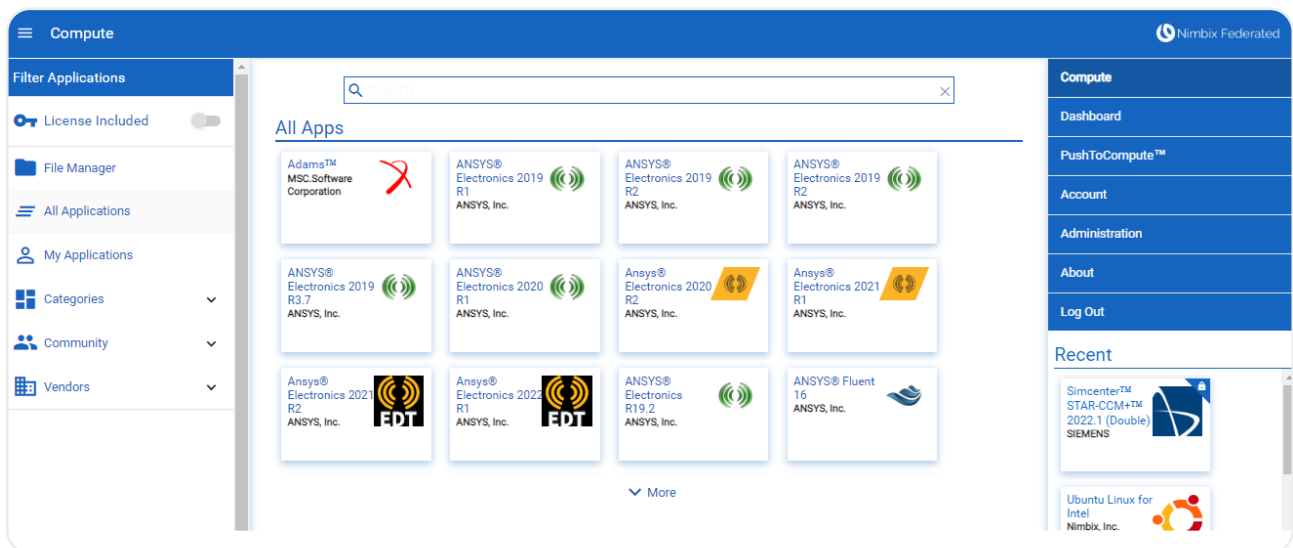
In response to this situation, Eviden launched the Nimbix Supercomputing Suite in 2022 - a set of flexible, secure high-performance computing (HPC) as-a-service solutions. This as-a-service model for HPC, AI and quantum in the cloud provides access to one of the broadest HPC and supercomputing portfolios - from hardware to bare metal-as-a-service to the democratization of advanced computing in the cloud across public and private data centers. It offers elastic, dedicated and federated supercomputing consumption models, but for the purposes of this publication, we will focus on federated only.

Nimbix Federated, a key pillar of the Nimbix Supercomputing Suite, helps solve the problems of cost, complexity and performance. For end users, it represents the state-of-the-art in responsive user interfaces, allowing point and click access from any device, any time, on any network. The rich catalog of ready-to-run workflows represents the most popular applications and vendors serving the engineering/simulation, life sciences and data science/AI spaces. These are the same applications users are already accustomed to running on their workstations or small clusters, tuned for maximum efficiency and delivered globally from the HyperHub™ application marketplace.

For business users, the JARVICE platform that powers Nimbix Federated delivers granular accounting, cost-controls and project management tools.

And finally, for operators, JARVICE allows easy integration into the global Nimbix Federated control plane, which can be located entirely within a specific geography (e.g. the European Union), and accessible ubiquitously from the public cloud. Centers looking to offer capacity via Nimbix Federated control the pricing, resource limits and "shaping," and enjoy automatic monetization any time users consume their systems. Eviden also supports the option of a private federation with its own restricted control plane, for groups of operators looking to further control access (e.g. within specific domains or countries).

Operators can choose to deploy Kubernetes on compute nodes for maximum isolation and flexibility, or leverage their existing Slurm and Singularity deployments without any additional software configuration management overhead. JARVICE provides a consistent user experience across virtually any containerized infrastructure and platform.



General Architecture and Operations

Nimbix Federated was designed from the ground up for multi-cluster, multi-cloud and multi-tenant operations. The underlying platform has powered Nimbix Elastic (“The Nimbix Cloud”) for the past decade, facilitating millions of large-scale jobs for thousands of users in nearly 70 countries around the world. The control plane is a service-oriented architecture (SOA) divided into two main parts: upstream and downstream.

Upstream (single)	Downstream (multiple)
<ul style="list-style-type: none">• Web-based user interface• Public API• Business layer• High-level scheduling• License-based queuing• SSO and identity management	<ul style="list-style-type: none">• Cluster-specific scheduling• Persistent storage• Compute

All control plane components - including the downstream agent - run as services on Kubernetes. This ensures global deployment capabilities (on any cloud or on-premises infrastructure), as well as fault tolerance and high availability. Operators choosing to interface with existing Slurm clusters simply need the ability to access the login node via SSH (and authorized private keys), using a single “service” user account. Two models are supported: one with a Kubernetes cluster in proximity to Slurm, and the other a “direct-to-Slurm” mode that can be driven directly from the cloud. For the former, operators must expose HTTPS to the upstream control plane, while for the latter, SSH. Similarly, when running compute directly on Kubernetes, an HTTPS port must be exposed to the control plan. This service should be secured via SSL-terminated ingress and basic HTTP authentication.

Architectural Fault Tolerance Capabilities

Nimbix Federated employs a job status reconciliation technology that can properly account for jobs running and terminating - even when remote downstream clusters are temporarily offline or unreachable. Likewise, a downstream cluster that fails to communicate with the control plane will not interrupt running jobs. The reconciliation process runs constantly and automatically synchronizes state whenever connectivity is available. The control plane itself is also almost entirely stateless, allowing scale-out of services for both capacity and load balancing.

Computational Accelerators

Nimbix Federated supports computational accelerators such as GPUs, FPGAs, and IPU, and when using Kubernetes employs a “best-fit” scheduling mechanism to ensure the scarcest resources are consumed last. For example, it will not place CPU-only jobs on GPU-capable nodes unless all CPU-only nodes are fully occupied.

The platform can also leverage GPUs for offloading 3D rendering for OpenGL-based applications.

Accounting and Billing

The Nimbix Federated service maintains per-user, per-job, per-project and per-tenant billing down to the second — and automatically bills tenants for monthly usage. Cluster providers are then paid accordingly. Billing reports by zone are available upon request, with constantly evolving self-service capabilities.



Remote Access

End users connect to Nimbix Federated over any network that can reach it. The control plane may reside in a central, continental location, while individual compute clusters may be spread across a large geography. For visualization jobs, users connect directly to the downstream(s) hosting them from their browser. The same HTTP(s) port that exposes the downstream API can proxy browser-based remote display sessions securely and with minimal latency.

For example, an end-user residing in Finland, connecting to the global Nimbix Federated control plane in central Europe, does not need to proxy via the control plane to access sessions remotely on a provider's downstream cluster in Finland. Users with access to multiple federated endpoints can further optimize their latency by selecting the cluster geographically nearest to them for compute and storage. The platform user interface provides a simple drop-down menu to change between federated zones.

Commercial Software Licensing

Nimbix Federated provides the ability to assign license servers to tenants and users. These license servers must be accessible from compute clusters, and enabling connectivity is part of the tenant onboarding process. Some applications support on-demand licensing and do not require deployment of persistent license servers. Naturally, for free and open-source software, license servers are not a concern. ISVs are free to provide their own click-through end-user license agreements (EULAs) in the web portal when users run applications from HyperHub™ application marketplace.


The platform also supports advanced license-based queuing, with per-project maximums and floors, and fair-share scheduling of license checkouts. It has the ability to pause solvers running for lower priority projects to allow higher priorities to run.

Storage Patterns

When using Kubernetes clusters for compute, Nimbix Federated relies on persistent volume claims (PVCs) that can refer either to statically or dynamically provisioned storage. Examples of statically provisioned storage include large shared or parallel file systems, organized hierarchically by project or user. Dynamically provisioned volumes can be either file or block storage, and work with all cloud infrastructure mechanisms. For on-premises use,

Nimbix Federated supports any storage topology that can be attached to Linux. When deploying a multi-tenant federated downstream, it can be configured with policies to automatically isolate tenant data in specific hierarchies on shared volumes. Tenants can decide whether to share datasets amongst their teams or provide further isolation on a per-user basis. When using Slurm, the platform supports file storage organized hierarchically by directory, and binds storage paths privately per tenant and per user. The same policies can be applied as those supported in the Kubernetes version, for sharing datasets or completely isolating users.

Data can either be transferred directly to the storage medium using provider-specific mechanisms, or by using the platform's built-in file manager application, which allows drag-and-drop for file upload and download. The advantage of file manager is that it supports single-sign-on at the control plane level, translating to a service account and isolated storage location on the target infrastructure. This eliminates the need to manage users downstream.



Security

Container Model

Nimbix Federated employs containerization to secure application workloads and isolate resources. All HyperHub™ containers, as well as self-service containers that tenants deploy in their own accounts, are in the open container initiative (OCI) format. ISVs and end users can easily build these containers on their laptops using Docker, or can leverage the built-in CI/CD pipeline with webhook capabilities that the platform provides for seamless container building and deployment. It supports any v1 or v2 registry and securely transmits “pull secrets” to compute targets ephemerally for running jobs.

On OCI-based systems such as Kubernetes, Nimbix Federated defines a v2 application model that prevents root access from inside containers and isolates networks by tenant. All persistent storage is provisioned and attached at the infrastructure level, preventing users from mounting their own storage or interfering with that of others.

On systems with Slurm, Nimbix Federated automatically converts and caches OCI containers to SIF and runs them with the same level of security that non-HyperHub™ Singularity containers would enjoy. However, rather than allow end users access to the underlying cluster via SSH in order to run MPI jobs, the platform orchestrates SSH services inside containers so that users cannot break out and run programs and scripts on the hosts. This is especially important in the federated model, since end users are essentially anonymous downstream, with both authentication and authorization (via SSO) taking place upstream. Preventing users from logging into the cluster directly allows operators to manage a single service account without additional exposure, greatly simplifying system administration.

For both Kubernetes and Slurm, operators can shape cluster resource selection by node label (Kubernetes) or partition (Slurm). This allows further isolation between research users native to the cluster and federated users coming in via an upstream control plane.

Multi-tenant Isolation

When using Kubernetes as the downstream cluster, the platform has the ability to ensure that no two tenants ever run containers on the same node(s) at the same time. While containers are fully isolated, this reduces attack vectors even further. Additionally, network traffic is automatically isolated between tenant jobs via software firewalling.

When using Slurm, operators may decide to offer nodeexclusive capabilities to tenants to avoid multiple containers running on the same node, thus achieving a similar result - albeit with less granularity.

Encrypted Transfers

The Nimbix Federated service encrypts all data it transfers — including for job control — via HTTPS. The best practice is to use CA-signed certificates and secure ingress with strong SSL encryption.

Even though URLs are transmitted securely via TLS in the HTTPS protocol, tenants also have the self-service option to eliminate random passwords from URLs for accessing remote jobs.


Single sign-on

Nimbix Federated supports external identity providers such as SAML2, LDAP and Active Directory. Providers are connected at the tenant level and the configuration is self-service for tenants. Identity brokers may also be used, provided they can produce SAML2 assertions.

When using LDAP or Active Directory, further authorization may be performed using LDAP substrings for specific resources (such as group or OU membership, etc.)

Resource and Application Limits

Tenants have the self-service ability to restrict not just cost but specific resource usage, as well as to limit the HyperHub™ catalog portions that users may access. Team defaults as well as per-user overrides are available. This allows for the protection of sensitive information in containers, such as reference data, from unauthorized users within a given organization. Limiting resource types can also prevent certain users from consuming capabilities they shouldn't be accessing and ensure appropriate availability of scarce resources such as GPUs.



Conclusion



HPC will continue to be the innovation engine for scientific and industrial breakthroughs. On our journey towards exascale, it is critical to achieve new heights in performance with unrivaled system efficiency, while democratizing access to these valuable HPC resources to wider communities.

If you are interested in getting involved with a supercomputing federation, either as a cluster operator or an industry user, click below for more information.

[Learn more](#)

Connect with us



eviden.com

Eviden is a registered trademark © Copyright 2023, Eviden SAS – All rights reserved.