

WHAT IS NODEWEAVER?

NodeWeaver is an innovative **hyperconverged infrastructure** – that integrates storage, networking and virtualization in a single system. It is built using the same principles of large scale systems used by Google and Amazon, making them available to small and medium enterprises. It has an easy to use interface and a management system that automates most tasks and simplifies activities that would otherwise require highly skilled expensive personnel.

NodeWeaver delivers linear and predictable scale-out without large up-front investments. Start small and expand your infrastructure as you need, one node at a time.

Having a unified compute, storage and networking platform enables your organization to deploy the applications you require without complex planning and procedures. NodeWeaver's Distributed File System aggregates internal and direct-attached storage resources across all nodes, presenting it as a single storage entity and making it available to all hosts. With NodeWeaver there is no need for external SANs as it's already included in the cluster itself.

■ Plug-and-play deployment

NodeWeaver's appliances provide you with a near plug-and-play platform. No more LUNs to configure on a SAN, planning of complex network configurations with redundant switches and paths, or workloads to adapt to a virtualization infrastructure. Just configure an IP address to reach the web interface and your infrastructure is ready to deliver services.

■ Scale out one node at a time

A platform that grows with your organisation. Each additional node adds resources to the overall capacity of the whole

system. While more CPUs, memory, storage and IOPS become available our intelligent allocation agent "weaves" resources together as a single entity and manages most of the tasks automatically to provide you with the best level of performance and resiliency achievable.

■ One interface to control everything

NodeWeaver provides a simple, web-based interface that can be changed easily to suit the needs of experienced administrators or end users. Decide what capability should be presented, and enjoy a single pane of glass to control everything from one page.

MAIN PILLARS

■ Distributed Computing

The NodeWeaver architecture is similar to Google's and Amazon's architectures in that it is a scale-out compute and storage infrastructure that eliminates the need for network storage. At the same time, NodeWeaver builds upon the latest research on distributed architecture and provides a general-purpose solution for virtualized environments. In addition to its scale-out capabilities, it has the same or better enterprise-class data management features that are commonly provided by advanced network storage solutions, including high availability, backup, snapshots, and disaster recovery.

■ Sea of Storage

Traditional virtualization still requires thinking in terms of physical units, LUNs, and managing each object as a separate entity. While it is the most obvious and the standard way of managing, it is not the best way.

Data flows like water, and by taking advantage of this idea we built NodeWeaver as a "sea of data", where each new node brings with it some capacity of holding the data. In the event of a component failure, data is free to flow to where it can be contained and managed properly. It flows on its own and for the first time, the user is free to care about what is really important, and not whether a disk has enough space for a virtual machine image or where to store backups. As it flows freely, data is adapting and moving to where it is needed most; there is no need for DRS or other technologies to optimally allocate resources, the resources will flow to where they are needed most. To help the user in expressing their requirements, the NodeWeaver sea of data is partitioned in dynamically allocated datastores, that provide a way to express preferences for use and placement of data as well as the number of replicas of every single virtual disk. Replicas and data integrity are constantly checked and managed without any need of human intervention.

Virtualization

The NodeWeaver architecture was designed from scratch to manage virtual machines - so it supports all hypervisor functions that are supported by the traditional network-storage based architecture, including live VM migration and high availability. Since in NodeWeaver the central element are Virtual Machines, it was designed to let the administrator forget about the complexity of traditional virtualization systems. Weaving together computing and storage, NodeWeaver bypasses performance bottlenecks caused by the division between the levels of computing and storage allowing it department to focus on the services that are delivered through the VM and not losing time on the service that delivers the VM.

SSD Tiering

NodeWeaver's architecture is designed to take advantage of enterprise-class solid-state disk drives (SSD). In hybrid configurations, where rotational drives coexist with solid-state ones, the tiering system automatically moves the most frequently used storage blocks on SSD drives. This intelligent placing maximizes the virtual disk's "hot zone" performance and prevents archive data, which is rarely accessed, from wasting precious solid-state storage. In this way performance and efficiency are maximized allowing substantial savings that can be measured in lower € / TB and € / IOPS ratios at the infrastructure level. The user can let NodeWeaver take care of the data positioning on the different types of media to ensure the best performance across the entire datastore, or, for special requirements, can decide to force certain VMs or individual images exclusively on SSD or rotational disks.

TECHNICAL HIGHLIGHTS

Full data integrity: NodeWeaver's Distributed File System automatically detects and fixes data integrity issues due to hardware failures, preventing silent data corruption. Thanks to the independent, per-block checksums, the data for each block is verified both during read and write processes and after it has been written to disk. Additionally, a background process periodically scans each block and verifies its checksum to prevent even single-bit alterations, and executes a self-healing and transparent repair through the additional replicas. This combination provides a long-term guarantee of perfect data availability even in the harshest environmental conditions.

Virtualization extensions: NodeWeaver implements most of the Hyper-V extensions designed to virtualize modern Windows systems in a more efficient way. In addition, the NodeWeaver hypervisor supports and correctly emulates a wide range of hardware devices, like the VGA card used by VMware and popular Intel ethernet cards.

RAIDless storage: NodeWeaver treats each disk unit as an independent system, with no need to spread writes to more than one disk like traditional RAID systems. This means that our platform has no inherent write penalty, and it does not suffer from substantial performance degradation when one disk is unavailable or when a new device is added to the infrastructure.

High efficiency SSD tiering: NodeWeaver appliances integrate both traditional rotational disks and enterprise-class SSD, and manages them together as an intelligent hybrid device. The SSD is used as a sophisticated cache device, transparent to the user, and thanks to its peculiar design multiplies the speed obtained by traditional caches especially when used for VDI.

I/O optimization: Each NodeWeaver appliance includes a special software subsystem that continuously collects data

on available resources, utilization and pending requests, and based on that data (along with that of the other nodes) continuously decides how to allocate CPU, I/O, and memory to guarantee a fair environment for each virtual machine. It also provides a continuous monitoring of how blocks are used and changed from the VMs, and moves them near (ideally in the same node) the hypervisor that is using them in order to reduce network contention and increase speed thanks to the local SSD.

Replace-on-write snapshots: Our distributed filesystem considers snapshots as a combination of online metadata and the sum of all the chunks that make up the online images that are visible in the management console. When the user (or the automated backup process) requests a snapshot to be taken, NodeWeaver copies only the metadata and creates a new snapshot that is indistinguishable from a file (no more treating snapshots as second class objects), completely read/write, and requires no additional space on disk, unless changes are committed. When a change is committed by a VM, the file system allocates a new group of chunks to hold just the differences, and this means that you can take snapshots of snapshots of snapshots... without any performance loss, and without the traditional "chain fragility" of traditional systems. You can delete, change or make a new snapshot of each of the intermediate files, with the guarantee that the system will always remain coherent. Thanks to this infrastructure, snapshots are atomic (instantaneous).

Integrated VDI/DaaS: Thanks to the SPICE protocol, NodeWeaver already provides all the necessary components for a virtual desktop deployment, both using imported physical-to-virtual machines and Virtual Desktops created from scratch. NodeWeaver, thanks to its integrated snapshotting, creates thinly provisioned desktops from one or multiple gold masters, and allows users to enjoy their virtual desktops through a selection of protocols, such as SPICE, RDP, VNC and NoMachine!, to fit any need.