



Power and
control at your
fingertips

Enter the new world of utility computing



Benefits

Extility is a utility computing platform, allowing hosting providers to offer their customers virtual dedicated servers through a self-service portal.

Offer a flexible service to your customers

- Bill per month, or bill per hour: it's your choice
- Allow your customers to self-provision their servers through a web-based portal
- Your customer can choose from a wide variety of operating systems and images and scale its system rapidly to meet demand

Offer your customers additional functionality

- Servers are fully manageable by the customer without intervention from your support department
- All storage is on high availability RAID systems
- Offer features not available on conventional dedicated servers, such as instant backups and appliance images

Instant provisioning

- Gain a competitive edge
- Generate earlier cash flow

Reduce capital costs

- Increase utilisation of compute resources
- Buy fewer larger identical units

Reduce power costs

- Server consolidation gives lower power cost per CPU cycle
- Reduce power use on idle machines

Reduce support costs

- No manual intervention in provisioning
- Reduce inventory management overhead
- Offer differentiated products on homogenous hardware
- Retain your existing managed hosting stack
- Low platform management requirements
- Automated recovery from hardware failure

Consolidate existing servers onto the platform

- The above benefits are available for internal servers too
- Ideal for test and lab servers
- Move existing customers or internal servers to Extility using P2V technology
- Ideal to simplify integration after acquisitions

Vendor independence

- Leverage your existing vendor relationships to provide the hardware for the platform.
- Use any compute node vendor
- Use any OS vendor
- Use any appliance image vendor
- Use an approved storage vendor

Features

Customised deployment & integration

- We will white-label the system so it looks like your own
- Multiple options for billing system integration (metered data level, invoice-line level, invoice level), or use our billing and collection module
- All actions can be driven through an API or a web interface

Provisioning

- Self-service via Control Panel or API
- Provisioning of Virtual Dedicated Servers
- Start, stop and delete Virtual Dedicated Servers
- Resize memory, storage and grow disks
- Remote KVM console available to customer via the web interface

Scalability

- Hierarchical architecture: tens of servers per node, up to 1,000 nodes per cluster, multiple clusters per controller
- A single cluster can serve tens of thousands of customers
- Multiple storage clusters (576TB per cluster) in storage subsystem

Security

- Virtual Machines provide VM isolation
- Xen + VT support gives complete logical and security separation between each guest and its host, and thus between customers
- System provides resource constraint at a per VM level - no 'root escalation' issues
- No credit card details on system (stays in billing system)

Network, VPN, firewall and load balancer options

- Central IP allocation function
- Multiple virtual NICs per virtual server
- Controlled provisioning of central firewall
- VPN & firewall software on customer OS
- Golden images of 'appliance' software creating virtual firewall, VPN, load balancer appliances plus private VLANs

Disks, clone, snapshot backup and restore

- Multiple disks can be attached to each server
- Each disk thin-provisioned to minimise storage use
- Instant clone, snapshot backup and restore
- Customers can generate their own images
- Combine, or use in addition to conventional backup

Virtual Data Centers

- Customers can organise their servers into virtual data centers
- Aggregated measurement and billing at the VDC level

API Access

- Customer API to allow for programmatic provisioning
- Licensee APIs for billing and node management

Commercial Model

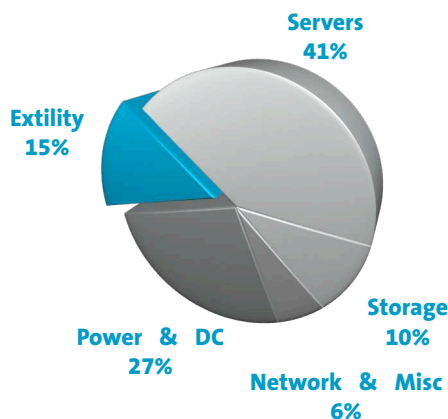
Specification and installation

We will mould our deployment to suit your needs. The first step is the recommendation of and agreement upon the hardware specification of the initial infrastructure deployment. This requires selection of hardware vendors, and determining the initial deployment size and likely growth rate, as well as specifying the integration requirements, particularly for the user interface and billing system. You are then responsible for procuring the hardware, and performing the physical set up. Following this, our engineers will install the software, integrate the systems, test the deployment, train your third line support staff, and hand the system over to you. Of course after handover we will continue to provide ongoing support as part of the package.

Pricing

Our pricing consists of two main elements: a set-up charge, consisting of a fixed fee plus an hourly charge for the integration work specific to your site, and then a monthly recurring fee which scales according to your use of the platform. We don't charge a large up front cost; rather, our revenue grows with the revenue you put on the platform, and the savings you make.

Typically, our charges will represent about 15% of your deployment costs (including amortised capital expenditure). So only a small gain in efficiency, such as reduced cost of servers or power is sufficient to make an investment in Extility pay off. Of course our virtualization platform will improve hardware utilization rates far more than that: utilization rates in data-centers are typically around 30% and may be as low as 10%; lifting utilization from 30% to between 60% and 90% will allow two to three times as much revenue to flow from the same capital expenditure, without a proportionate increase in your power bill. Plus your customers will appreciate the extra flexibility you can offer them, and you will appreciate the earlier revenue stream and reduced management overhead.



Typical deployment cost breakdown

Technical Specifications

Architecture

Extility is built using an inherently scalable architecture. The main components are: a user interface and control system, one or more XVP clusters, a storage subsystem made up of several redundant storage nodes, and a network subsystem.

Nodes and control system

The user interface and control system provides the portal for end-users and operations staff, implements the business processes and operates the link to the billing system. It consists of a redundant pair of XVP controllers which can control multiple clusters. Each cluster supports up to about 16,000 cores, but a lower limit per cluster may be desirable. An XVP cluster is made up of many XVP nodes, controlled by a redundant pair of XVP managers; these are combined with the XVP controller in single cluster configuration. Each XVP node is a high specification diskless multi-core machine running customer virtual machines. We recommend using nodes with at least 16 cores. Each node runs the Xen hypervisor (though the system is designed to support multiple hypervisor vendors), which securely partitions the node at a logical level into virtual machines: one host and multiple guests. The host operating system on each node is Linux, running our node control software. The guest virtual machines run the end-user operating systems, and are instantiated by the controller via the node control software.

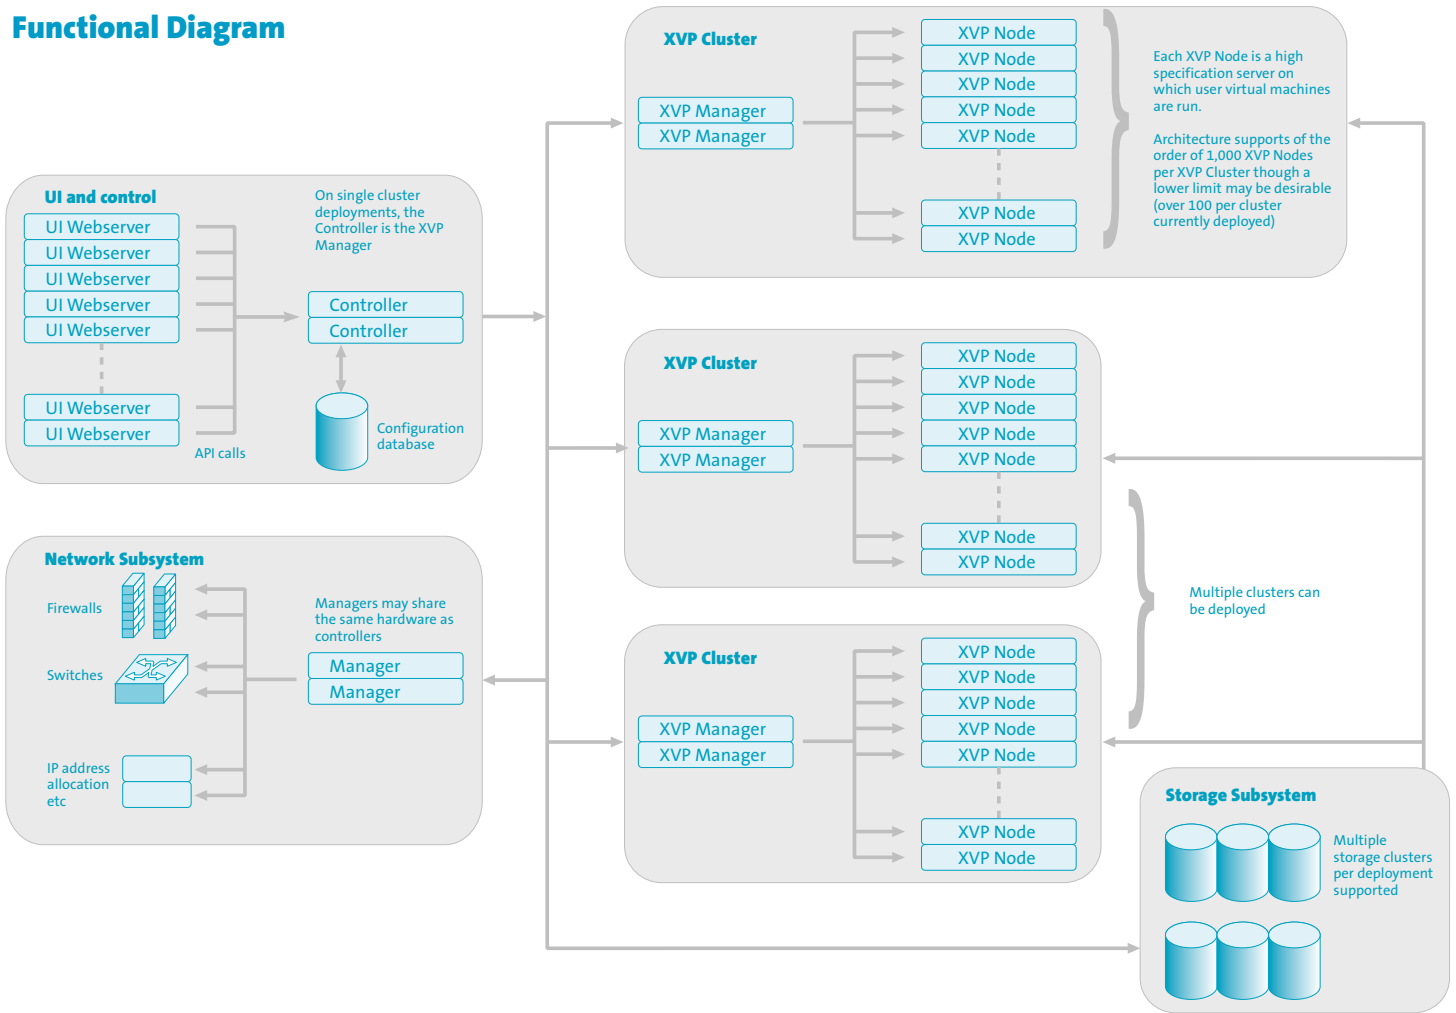
Storage

The storage subsystem consists of one or more Sun 7000 series high availability clusters (though we hope to support other vendors in future). These provide iSCSI LUNs to the end-user operating systems, looking just like physical disks. However, the storage system allows for instant clone and restore, as well as thin-provisioning, giving huge savings on required storage in addition to maximizing reliability.

Network

The network subsystem configures the layer two switched architecture and external firewalls. It also provides IP address management, tracks bandwidth, and provides reporting and billing. A typical configuration consists of switched 802.1Q VLANs, trunked to the XVP nodes. Three physically separate networks (A, B, C) carry customer traffic (which itself may consist of multiple VLANs), storage traffic, and management. Each XVP node has six gigabit ethernet ports spread across three gigabit ethernet cards, with the networks split so that failure of any card still leaves a connection to each network. Only appropriate VLANs (public network plus the customer's own private VLANs) are visible to guest virtual machines. Two top-of-rack switches per rack aggregate traffic on a 10 Gig-E capable network, to which the storage subsystem is directly attached.

Functional Diagram



Network Diagram

