Iron Networks Microsoft Fast Track Architecture

Reference Architecture Technical White Paper

IronPOD Product Family

Prepared by

Iron Networks Engineering
Robert Keith

Revision 1.0.9
Published April 2014

http://www.ironnetworks.com

Contents

1 Introduction

CIOs and IT decision makers have many pain points such as staffing skills shortages, managing multiple vendors, scaling out service infrastructures, keeping up with new and quickly evolving technologies while still managing risk and increasing business agility by increasing the services provided while still reducing operating costs.

Iron Networks with Microsoft's guidance developed the IronPOD family of products which are Microsoft Fast Track certified and conforms to the Microsoft Cloud Reference Architecture. This architecture leverages Microsoft's broad experience providing services and technologies to various industries, and providing cloud solutions such as Office 365 and Azure. This architecture design details the conceptual, logical and physical attributes for the Infrastructure-as-a-Service Cloud Service model and provides solutions for many of the CIO's pain points.

This document is a functional specification which details the Iron Networks IronPOD family of converged infrastructure solutions that combine servers, storage, networking and infrastructure management into an integrated, tested and optimized public, private and hybrid cloud infrastructure. These solutions include multi-tenancy, fabric management, compute, network and storage fabric components with automated deployment technologies.

1.1 Iron Networks

Iron Networks is a hardware product and services company and a Microsoft Fast Track partner with a diverse customer base and product portfolio ranging from low energy micro servers to mission critical converged data center and cloud platforms.

Iron Networks has been building Microsoft based hardware appliances since 2007. Iron Networks is also well positioned as an OEM partner with many Tier-1 hardware manufacturers as well as the industries' largest commodity hardware manufacturers. The company provides an overarching set of services such as manufacturing, quality testing and integration and supply-chain worldwide combined with 24/7 support staffs, professional services and call-home managed services. This combination of services provides a complete cloud product and service solution.

Iron Networks products include security appliances, remote access solutions, Big Data and Hadoop products, storage products, networking and full-scale multi-rack data center solutions.

This agnostic hardware position allows Iron Networks to provide an integrated solution utilizing the "best of breed" components from several manufacturers to provide customers a Fast Track certified solution based on their own hardware standards and preferences.

The IronPOD family is top-to-bottom a software defined architecture based on Microsoft Windows 2012 and System Center 2012 SP1. The product solutions include storage hardware designs based on Windows Storage Spaces as well as Virtual Networking technologies based on Hyper-V Virtual Networking.

This software defined architecture provides the latest evolution of hardware and software technologies which provides complete customer flexibility and scalability while providing the greatest performance at the best achievable market prices and lowest TCO.

1.2 Industry and Technology Trends

The technology landscape of IT changes rapidly with constant and ongoing technology developments. CIOs and decision makers must stay abreast of technology trends make plans and budgets accordingly. To remain relevant and provide the best set of services requires adapting new products and services to meet marketplace and user demands.

Current technology trends include:

- Cloud Computing
- Fabric Management
- Software Defined Everything
- Solutions based on Commodity Hardware
- Virtualization
- Network Protocol Convergence
- Big Data
- Remote Computing and Mobile Devices

1.2.1 Cloud Computing

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable resources (such as networks, servers, storage, applications and services) that can be rapidly provisions and released with minimal management effort or service provider interaction.

NIST defines cloud computing as:
"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (such as networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models. " [Excerpt from NIST Definition of Cloud Computing (Mell and Grance 2011)]

Essential Characteristics:
On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (such as mobile phones, tablets, laptops, and workstations).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (such as country, state, or data center). Examples of resources include storage, processing, memory, and network bandwidth.

Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (such as storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models:

Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (web-based email) or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (host firewalls).

Deployment Models:
Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (such as mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Public cloud. The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (such as cloud bursting for load balancing between clouds).

1.2.2 Fabric Management

The fabric architecture is comprised of two parts. The first part is the fabric, which is the physical infrastructure comprised of servers, storage, and the network that will host and run all customer or consumer virtual machines. The fabric is the cloud infrastructure. The second part is fabric management, which is the centralized control, monitoring and management of the cloud infrastructure.

In the Microsoft Cloud Architecture, System Center provides the foundation for fabric management. SCVMM and SCOM are the foundational modules which manage the virtualization infrastructure and monitor for health. The additional modules of System Center provide higher level functions for management of the fabric.

1.2.3 Software Defined Everything

The term Software-Defined Everything (SDE) includes software-defined networking (SDN), software-defined computing, software-defined data centers (SDDC), software-defined storage (SDS) and software-defined storage networks.

With Software-Defined Everything, the entire infrastructure is virtualized and delivered as a service. In a Software-Defined Everything environment, management and control of the networking, storage and/or data center infrastructure is automated by fabric management software running on generic servers rather than by specific physical hardware components.

Microsoft has extensive experience providing services and technologies to various industries, and providing cloud solutions such as Office 365 and Azure. The Microsoft Cloud Architecture leverages this experience.

Microsoft is the world's largest software company and is well positioned to provide the technology foundation for a Software Defined Everything infrastructure.

1.2.4 Commodity Hardware

Technologies such as Virtualization and Software Defined Everything have made specific purpose and expensive hardware less of (or not) a requirement.   In server farms, when a server fails, the work loads are running in virtual environments which migrate almost immediately to other available servers.

Windows Storage Spaces provides highly sophisticated, software defined, storage solutions using standard servers with JBOD storage systems attached.   Complex and expensive SAN and NAS solutions are no longer required to provide the advanced features, capacity and performance required for large computing workloads.

Technologies such as OpenFlow allow the management and complexity of the networking fabric to be centralized to a fabric management system while the actual forwarding of network packets is handled by simpler network switches which are optimized for capacity and reduced latency.

1.2.5 Virtualization

Virtual machines, virtual storage and virtual networking technologies have abstracted workload computing from the actual physical hardware.   This increases flexibility since the workload applications are no longer locked to any specific hardware system. High availability software systems and clustering allows the hardware infrastructure to be slightly over provisioned, where when hardware outages occur, the workload is redistributed automatically and the loss in capacity is absorbed by the extra hardware components.

This flexibility and increased reliability options comes with an increase in complexity as well. Virtualization adds a layer of software between the physical hardware and the operating systems running the workload applications.

Fabric Management systems manage the operation of virtualization and isolates the administrator from the complexity of the underlying hardware and software abstraction layers.

1.2.6 Network Protocol Convergence

Network technologies are increasingly expanding and growing more complex.   Today's switches and routers manage over 6000 network protocols.

Current trends are to move network traffic loads to packet forwarding hardware devices.   Policies which control the packet forwarding behavior and policies are managed via the Fabric Management systems. The hardware platform is optimized for performance, and the complexity of managing networks is handled in software on generic servers removed from the packet forwarding hardware.

Layer 2 switching has continued to increase in complexity, requiring ever more complex and expensive hardware systems. The trend is to move more traffic to layer 3 routing protocols, and away from the layer 2 switching protocols. This provides flexibility and additional features for scaling systems since the traffic is available to be routed throughout the entire infrastructure.

1.2.7 Big Data

Big Data and Data Analytics using systems such as Hadoop require new infrastructures and technology skills. These infrastructures require additional hardware infrastructures, and are typically not based on virtualized infrastructures.

Most businesses today are in process or planning to implement Big Data projects.

Iron Networks provides prebuilt and pre-integrated Big Data infrastructures. These of-the-shelf systems can be installed like a simple appliance and is available to use almost instantly. This provides a quick bootstrap to a Big Data project and allows the IT administration staff time to train and prepare to scale-out the Big Data hardware environment.

1.2.8 Remote Computing and Mobile Devices

Remote Access to corporate systems traditionally was provided by VPN technologies. These solutions have been complex to secure and manage, and difficult to scale. The current explosion of remote devices includes tablets and smart phones have expanded the management complexities.

Microsoft provides several technologies for remote access including DirectAccess, UAG and WAP. These technologies work together to connect and manage Active Directory domain managed remote endpoints as well as unmanaged endpoints like smart phones and tablet devices.

Iron Networks provides pre-integrated security and remote access hardware systems which extend the functionality and access of these remote devices as well as providing extended security. These additional functions include multi-tenancy remote access solutions for cloud providers.

2 Microsoft Fast Track Program

The Microsoft Fast Track Program is a joint reference architecture for building private clouds that combines Microsoft software, consolidated guidance, and validated configurations with OEM partner technology, including computing power and storage architectures and value-added software components.

Microsoft Private Cloud Fast Track solutions provide a turnkey approach for delivering preconfigured and validated implementations of the private cloud. With local control over data and operations, IT professionals can dynamically pool, allocate and mange resources for agile infrastructures as a service. In addition, business managers can deploy line-of-business applications with speed and consistency using self-provisioning automated data center services in a virtualized environment.

2.1 Business Benefits

The Microsoft Fast Track Program provides reference architecture for building private clouds on each organization's unique terms. Each fast-track solution helps organizations implement private clouds with increased ease and confidence. Among the benefits of the Microsoft Hyper-V Cloud Fast Track Program are faster deployment, reduced risk, and a lower cost of ownership.

The customer has the choice of building the solution by using the reference implementation guidance from Microsoft or purchasing a solution from a Microsoft hardware partner that couples the guidance with optimized hardware configurations. Although both options decrease the time, cost and risk in implementing private clouds, purchasing a reference implementation from Microsoft hardware partner will result in the fastest, lowest-risk solution. This is due to the fact that in this option, all the hardware and software best practice implementation choices have been determined by Microsoft and its hardware partners' engineering teams. As a result, this will also prove to be the most inexpensive option.

Faster-deployment:

- End-to-end architectural and deployment guidance
- Streamlined infrastructure planning due to predefined capacity
- Enhanced functionality and automation through deep knowledge of infrastructure
- Integrated management for virtual machine (VM) and infrastructure deployment
- Self-service portal for rapid and simplified provisioning of resources

Reduced risk:

- Tested, end-to-end interoperability of compute, storage and network
- Predefined, out-of-box solutions based on a common cloud architecture that has already been tested and validated
- High degree of service availability through automated load balancing

Lower cost of ownership:

- A cost-optimized platform and software-independent solution for rack system integration
- High performance and scalability with Windows Server 2012 operating system advanced platform editions of Hyper-V technology
- Minimized backup times and fulfilled recovery time objectives for each business critical environment

Technical Benefits

The Microsoft Private Cloud Fast Track Program integrates multiple Microsoft products and technologies, in addition to hardware requirements, to create reference implementation guidance. If the solution is purchased from a Microsoft hardware partner, the reference implementation guidance is implemented with partner hardware and sold as a reference implementation. Whether the customer decides to implement the Microsoft-validated reference implementation guidance with their own hardware or with hardware from a Microsoft partner, it goes through a validation process. In either case, Microsoft and its hardware partners have created a solution that is ready to meet customer needs.

3 Architecture Overview

3.1 Objectives

The architectural objectives of the IronPOD families is to provide a set of products for the Public and Private cloud based on the Microsoft Cloud Reference Architecture using Fast Track design principles.

Additionally, Iron Networks includes the following design goals

- Simple Deployment
- Simple Management
- Maximum Performance at the Lowest Cost

Simple Deployment

Iron Networks has developed an automated deployment technology which builds all of the IaaS components from bare-metal.

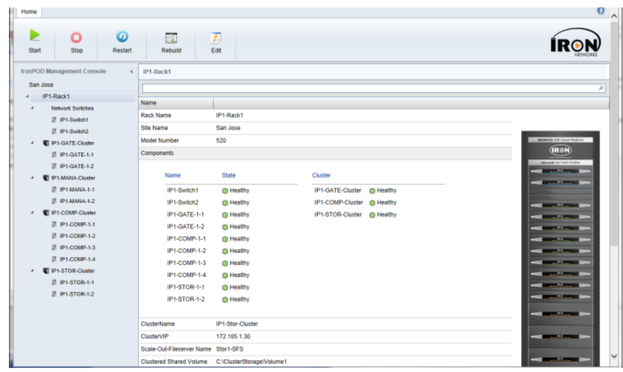

The installation and configuration of a system with hundreds of software and hardware components is labor intensive, complex and error-prone.

The Iron Networks Deployment system builds any module of IronPOD, or deploys racks of equipment simultaneously with a single click.

Simple Management

Cloud infrastructures are complex and constructed of many technology products. Each of these products has their own management interface technologies and underlying philosophies. Major technologies such as storage, networking and virtualization have dozens to hundreds of product manuals. The cost of maintaining a skilled staffs and training is enormous.

The IronPOD products are based on Microsoft Windows 2012 and Microsoft System Center 2012 R1. Software Defined Everything designs has allowed the convergence of technologies into Microsoft Windows. The management of the entire technology stack, from storage, networking to virtualization and including the operating system is a single technology infrastructure from a single vendor. IT administrators familiar with Windows will be comfortable with the entire Cloud infrastructure management. Technology convergence is driving the reduction of skills, training, risk and TCO.

Iron Networks integrates and tests all of the hardware components, then tests the entire product with the software configuration to guarantee the entire system functions without conflicts and with optimized performance. The IronPOD systems are shipped and supported as a single product. Iron

Networks provides 24/7 support and onsite support options, so support for Cloud IaaS systems are provided from a single support service.

Any issues with hardware, software, performance, supply chain or managed services are provided from a single source, which is the simplest possible management solution.

Maximum Performance at the Lowest Cost

Iron Networks is located in the heart of the Silicon Valley, and is an OEM manufacturer with strong relationships with most of the computer hardware companies.   The latest technologies on the market often have the best performance verses cost ratio.   Iron Networks' design goals take advantage of the newest products with the best performance available for the lowest costs.

4 Cloud Architecture

The Public Cloud is a virtualized Infrastructure-as-a-Service service provided to the public.

Virtualization technologies have enabled this service to be viable. Cloud design criteria such as resiliency, scalability, elasticity, multi-tenancy, fabric scale-out designs and management are important design goals for any modern infrastructure.

The IronPOD architecture provides the modules to build a flexible, scalable and less expensive solution to public cloud providers.

4.1 Elasticity

The cloud infrastructure should appear to customers as infinite capacity. In reality the administrator has to have the ability to scale up and scale out systems while live and in operation. To accomplish this requires several technologies including virtualization, storage and advanced networking protocols.

4.2 Elastic Storage

Microsoft Storage Spaces has a disk storage pool abstraction layer. Disks are grouped into storage pools. These pools can be defined for different purposes including workload requirements, performance, cost and physical locations or hardware enclosures. This software defined storage construct also allows disk drives to be added to the pool dynamically, allowing the capacity to be expanded (scaled out) while the pool is live and active.

4.2.1 Storage Tiering

Storage tiering is the practice of physically partitioning data into multiple distinct classes, such as price or performance. Data can be dynamically moved among classes in a tiered storage implementation based on access, activity, or other considerations.

SSDs can be added to pools, and with Storage Spaces Auto Tiering enabled, increases performance, scaling up the storage infrastructure.

The IronPOD storage products include configurations which enable Windows Storage Spaces Auto Tiering. This technology automatically moves "hot" storage blocks to the SSD drives, greatly increasing storage performance.

The IronPOD storage products also have product options which include configurations with varying disk types including SSD drives, 15000 RPM HDD drives and 7200 RPM drives. This provides many options to create storage performance tiers.

In the scenario described below, four storage performance tiers are defined:

- SSD Tier (Fastest)
- SSD/HDD Auto Tier
- 15000 RPM Tier
- 7200 RPM Tier (Slowest)

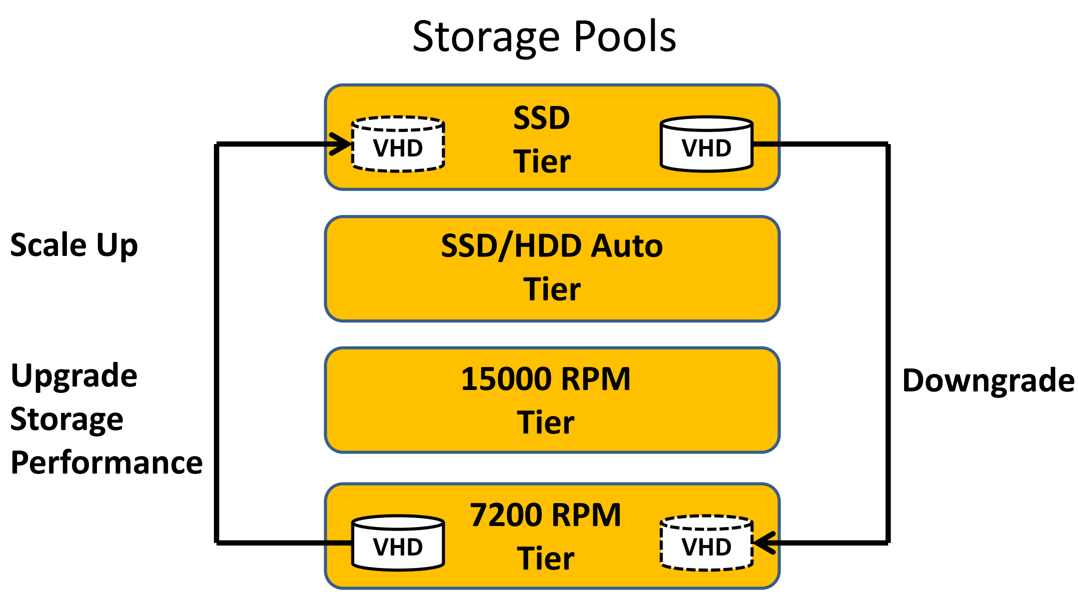

Upgrading the performance can be accomplished by migrating the customer Virtual Hard Disks (VHD) from a slower storage pool to a faster pool.

4.2.2 Live Migration of VHDs

Virtual Hard Disks (VHDs) can be moved live across storage devices and pools. The Virtual Machine (VM) will continue to run and does not have to be shut down. This allows the administrator to upgrade (scale up) or downgrade the performance of the customer's storage as desired or required, based on SLAs and customer requirements.

When correctly architected high speed network fabric technologies are in place and utilizing technologies such as RDMA from Microsoft Windows, Application and Server VMs can be distributed across a data center with very little performance degradation. For example, a server VM on rack 1 can access the OS VHD on rack 2 and the application database VHD on rack 3. This allows the administrator to manage storage infrastructures, to move VHDs around the storage fabric as desired, and add additional storage devices and racks (scale out), without the requirement to carefully place VHDs within the storage fabric.

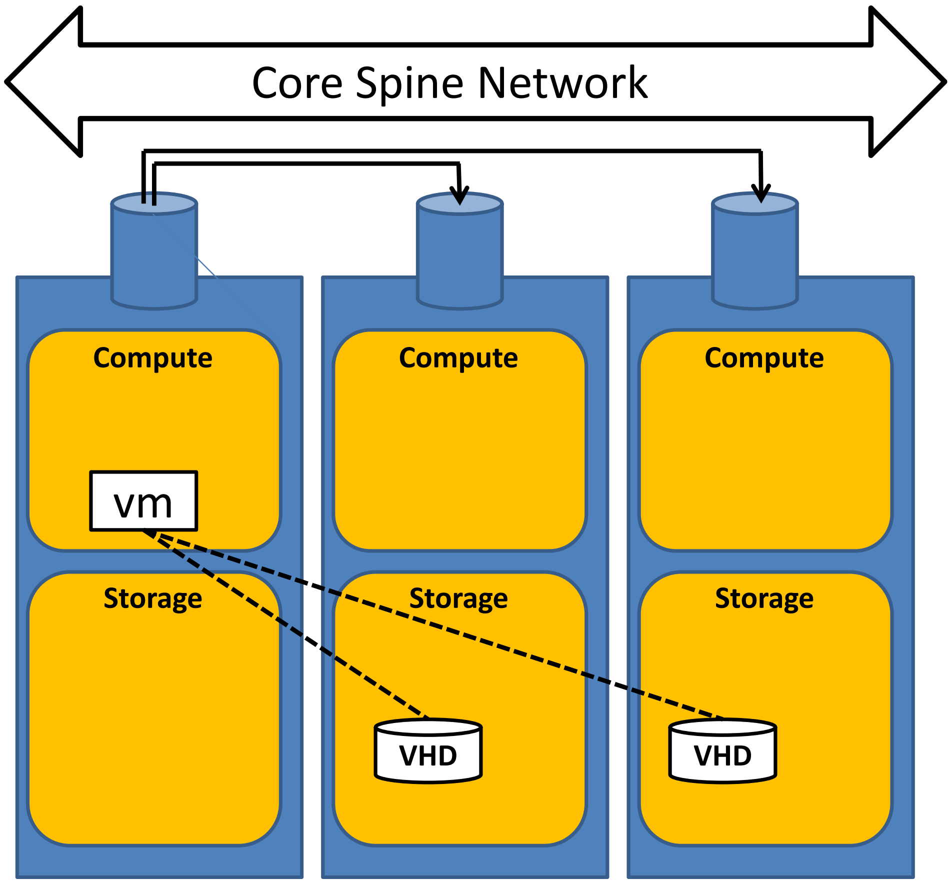

4.3 Multi-Tenancy

Multi-tenancy refers to a principle in which an infrastructure can be logically subdivided and provisioned to organizations or organizational units. The traditional example is a hosting company that provides servers to multiple customer organizations. Increasingly, this model is being used by organizations with a centralized IT department that provides services to multiple business or organizational units and treats each as a customer or tenant.

4.4 Resiliency

From the consumer's perspective, cloud services should always be available. The consumer should never experience an interruption of service, even if failures occur within the IaaS environment. To achieve this, a provider must have application resiliency and infrastructure redundancies in a highly automated environment. Much like the perception of infinite capacity, the perception of failure proof operation can only be achieved in conjunction with high availability design principles.

High ability design principles include:

- No Single Points of Failure
- Automatic Failover
- Application Resiliency Awareness
- Redundant Network Paths
- Redundant Components

No Single Points of Failure

A highly available design must have a “No single point of failure” criteria in every aspect, from every physical component through the fabric tiers and designed into every application.

Failure of a single critical component can disable an entire infrastructure if not using this design.

Automatic Failover

All failures should result in triggering an automated failover to an active available resource. No manual intervention should be required and the failure should not cause any interruption in service or loss of data.

Examples of automatic failover technologies in an IaaS infrastructure are numerous. These include Windows clusters for compute and storage node failover, teamed network ports, multiple ingress network paths with Layer 2 or Layer 3 failover and storage mirroring or parity technologies.

Excess capacity must be must be provisioned to be available to take the place of failed components.

Failures at worst should only reduce capacity over time and not affect availability.

Application Resiliency Awareness

Highly available applications must be designed with resiliency, including running on multiple server nodes, being cluster aware to maintain data integrity during failovers and storing data in a manner that is replicated.

In the case of failures, where automatic failovers occur, applications can become corrupt if they are not designed with this infrastructure awareness.

For example, applications such as Microsoft Exchange and Apache Hadoop provide data parity and replication at the application layer and providing redundancy at the storage layer is undesired. Applications like Microsoft Lync are not compatible with Hyper-V Replica and require special handling for IT operations.

Redundant Components and Network Paths

Inevitably hardware will fail, become damaged or be misconfigured. To guarantee no service interruptions, every component must be replicated and available.

Providing redundant paths from the server through all the network tiers to the provider's core network tier is critical for high availability and resiliency. Technologies like network adapter teaming or the spanning tree algorithm can be utilized to provide redundant path availability without looping.

Any correctly architected cloud infrastructure will have redundant network paths to every resource.

- Ingress network connections must include multiple connections to multiple provider core switches. Traffic should failover almost instantly and automatically.
- Every compute node should have redundant teamed network ports connected to multiple redundant switches.
- Every network switch should have all network traffic routed automatically to alternate redundant switches. Any critical switch that fails should not create any service interruption.
- Storage servers should be clustered with excess compute capacity where the loss of a storage server will cause no interruptions.
- The data paths from storage servers to the disk drives should be redundant and path failures automatically routed around with a technology such as MPIO.
- Any and all storage components should be mirrored or replicated using technologies such as storage spaces parity or mirroring, or an application equivalent.

4.5 Virtualization

Virtualization is provided at multiple layers, including storage, network, and server. Virtualization supports resource pooling at each of these layers and abstraction between the layers for greater efficiency.

Storage Virtualization

Storage virtualization uses software to create a Virtual Hard Disk (VHD) that emulates a physical disk drive. Storage virtualization refers to the abstraction (separation) of logical storage from physical storage so that it can be accessed without regard to physical storage or heterogeneous structure. This separation allows increased flexibility for how system administrators manage storage for end users.

Network Virtualization

Network virtualization uses software to create a Virtual Local Area Network or VLAN that emulates a physical local area network. Network virtualization is the process of combining hardware and software network resources and network functionality into a single, software-based administrative entity known as a virtual network.

Server Virtualization

Hardware virtualization uses software to create a virtual machine that emulates a physical computer. This virtualization creates a separate operating system environment that is logically isolated from the host server. By providing multiple virtual machines at once, this approach allows several operating systems to run simultaneously on a single physical computer.

Hyper-V technology is based on a 64-bit hypervisor-based microkernel architecture that enables standard services and resources to create, manage, and disable virtual machines. The Windows hypervisor runs directly above the hardware and provides strong isolation between the partitions by enforcing access policies for critical system resources such as memory and processors. The Windows hypervisor does not contain non-Microsoft device drivers or code, which minimizes its attack surface and provides a more secure architecture

4.6 Physical Network

The physical network as described is the physical switches, cables and network ports on each server and storage node.

The network is defined in access tiers

- Core Network
- Aggregate Network
- Tenant Network
- Storage Network

There are other networks such as the Out-of-Band Management network and the data path network from the storage servers and the disk drives. These networks are described elsewhere in this document.

4.6.1 Core Network

The core tier is the high-speed backbone for the network architecture. The core is typically comprised of two modular switch chassis that provide a variety of service and interface module options. This network tier may be called the “Spine” network, and is the interconnect between multiple equipment racks.

4.6.2 Aggregate Network

The aggregation (or distribution) tier consolidates connectivity from multiple core tier switches. This tier is commonly implemented in end-of-row switches, a centralized wiring closet, or a main distribution frame (MDF) room. The aggregation tier provides high-speed switching and more advanced features, such as Layer 3 routing and other policy-based networking capabilities. The aggregation tier should have redundant, high-speed uplinks to the core tier for high availability and manage the failover for uplink failures.

4.6.3 Tenant Network

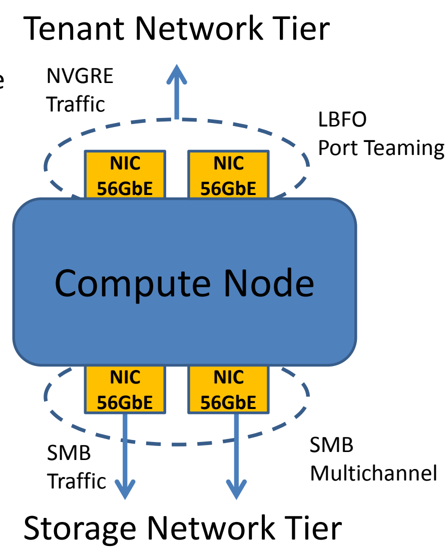

The tenant network resides between the aggregate network tier and the compute nodes. This tenant network fabric includes the multi-tenancy technologies including Network Virtualization and tenant isolation. The tenant network tier is where the NVGRE encapsulation is processed, and includes the Virtual Networking NVGRE gateway services.

4.6.4 Storage Network

The storage network resides between the compute nodes and the storage server controller nodes. This network fabric includes the storage technologies such as SMB3, RDMA and RoCE protocols.

4.7 Out-of-Band Management Network

Out-Of-Band (OOB) management is the ability to access all critical devices at the most primitive console level remotely, even when the main infrastructure network is unavailable and the device is powered down. For servers, this is access to the BIOS level console. For network switches the direct physical console.

OOB access should include all manageable devices including the Power Distribution Unit (PDU) hardware.

This OOB network is often called the Access network or the Management network.

OOB access must be provisioned by alternate physical network paths including dedicated switches and dedicated server ports.

4.8 Fabric Management

Fabric management is the concept of treating discrete capacity pools of servers, storage, and networks as a single fabric. The fabric is then subdivided into capacity clouds, or resource pools, which carry characteristics such as delegation of access and administration, service level agreements (SLAs), and cost metering. Fabric management centralizes and automates complex management functions that can be carried out in a highly standardized, repeatable fashion to increase availability and lower operational costs.

4.9 Process Automation and Orchestration

The orchestration layer that manages the automation and management components must be implemented as the interface between the IT organization and the infrastructure. Orchestration provides the bridge between IT business logic, such as "Deploy a new web-server virtual machine when capacity reaches 85 percent," and the dozens of steps in an automated workflow that are required to implement such a change.

Ideally, the orchestration layer provides a graphical interface that combines complex workflows with events and activities across multiple management system components and forms an end-to-end IT business process. The orchestration layer must provide the ability to design, test, implement, and monitor these IT workflows.

4.10 User Self-Service

Self-service capability is a characteristic of public and private cloud computing, and it must be present in any implementation. The intent is to permit users to approach a self-service capability and be presented with options that are available for provisioning. The capability may be basic (provision a virtual machine with a predefined configuration), more advanced (allow configuration options to the base configuration), or complex (implement a platform capability or service).

Self-service capability is a critical business driver that enables members of an organization to become more agile in responding to business needs with IT capabilities that align and conform to internal business and IT requirements.

The interface between IT and the business should be abstracted to a well-defined, simple, and approved set of service options. The options should be presented as a menu in a portal or be available from the command line. The business can select these services from the catalog, start the provisioning process, and be notified upon completion, at which point they are charged only for the services that are actually used.

5 Private Cloud Architecture

The design principles for private cloud providers are the same for public cloud architectures. Resiliency, scalability, elasticity, multi-tenancy, fabric scale-out designs and management are important design goals for any modern infrastructure.

When you take a service provider's approach for delivering information technology, a key capability is to be able to meter resource utilization and charge users for that usage. Historically, when IT departments have been asked to deliver a service to the business, they purchase the necessary components and then build an infrastructure that is specific to the service requirements. This process can result in an increase in time-to-market, higher costs because of duplicate infrastructures, and unmet business expectations of agility and cost reduction. Further compounding the issue, this model is often used when an existing service needs to be expanded or upgraded.

IT departments can transform their organization by taking a service provider's approach. When infrastructure is provided as a service, IT departments can use a shared resource model that enables economies of scale, and they can also combine other private cloud architecture principles and concepts to achieve greater agility for providing services.

The private cloud model provides much of the efficiency and agility of cloud computing in addition to the increased control and customization that is achieved through dedicated private resources. Many organizations require the control and flexibility required to provide the benefits of the private cloud as well as control their own destiny and guarantees for their internal security and SLA requirements.

Since most of the critical requirements of the private cloud are the same as for the public cloud, the solution set is similar with the possible exception of

- Commodity Hardware
- Service usage billing
- Self-Service provisioning
- Multi-tenancy

5.1 Commodity Hardware

Public hosting companies sell services at a price where profits are based on margins above cost and competitive advantages require the hosting company to have the lowest possible operating expenses. These companies typically operate thousands of servers and have larger IT administration staffs than corporate IT departments.

Since "white box" hardware solutions can have higher support footprints than enterprise server solutions, hosting cloud providers have larger homogenous server farms, and may be better leveraged to take advantage of lower hardware costs of commodity hardware.

Corporate IT departments are usually not profit centers and prefer the better support, quality and management infrastructures as well as the supply-chain advantages of tier-1 hardware providers.

Iron Networks addresses this situation with

- Iron Networks product offerings provide both commodity product components as well as solutions based on enterprise class hardware from well-known Tier-1 manufacturers. This provides both options for IT architects and decision makers and allows companies to maintain corporate IT hardware standards.
- Iron Networks performs a series of quality tests in their manufacturing process, then provides 24/7 enterprise support and supply chain options for both the enterprise components as well as the solutions based on commodity hardware. This allows companies to select hardware with lower TCO costs.

5.2 Service Usage Billing

Providing accurate cost models for cost accounting systems is critical, and cloud architectures must provide service-use data to be able to accurately measure the cost of operations and providing service to each tenant and bill appropriately.

Corporate datacenters often service multiple divisions and have a similar requirement to support charge-back accounting systems. This provides the basis for IT departments to act like a profit center calculating ROI's verses TCO, and increase user satisfaction.

5.3 Self-Service Provisioning

Public cloud providers have a critical requirement to allow end-customers to self-provision resources. While private providers have less of a requirement for self-provisioning, the larger the user base grows, the larger the IT administration staff costs rise and automating some or all of the administration and resource provisioning reduces operating burdens and increases the service levels and increases user satisfaction.

Resources available to self-provision include

- Generation of new VMs based on a workload template
- Provision additional storage resources
- User and group administration
- Generate new tenant virtual networks
- Provisioning of remote access
- Security administration
- Upgrade and downgrade VM and Storage VHD drives to faster or slower fabric resources

5.4 Multi-Tenancy

Public cloud providers have a core requirement to provide isolated and independent infrastructures to each customer where each customer has what looks like an infinitely scalable environment with his own VMs, Storage and IP address ranges. This fabric is provided by Microsoft's SCVMM and Hyper-V Virtual Networking technologies.

Private companies can benefit from this technology with

- Scalability – VM Management, Live Migrations,
- Self-Provisioning Management
- Corporate Acquisitions
- Test Environments

5.5 Scalability

Large data centers with a large number of active VMs run into a scaling issue when managing VLANs across multiple switches and racks. Multi-Tenant Hyper-V Virtual Networking removes the requirement for VLANs.

As well, the Hyper-V Virtual Networking allows live-migrations to cross subnets, so VMs and VHD resources can reside anywhere inside the datacenter fabric without the requirement to strategically place VMs close to the companion VHDs.

5.6 Self-Provisioning Management

Some data center resources have the possibility of being self-provisioned.

Allowing corporate users to self-provision resources reduces the administrator overhead and provides a higher level of service to the end customer and thereby increases user satisfaction.

5.7 Corporate Acquisitions

Corporate acquisitions often bring new "foreign" environments into the corporate infrastructure. With Multi-Tenant Virtual Networking technologies, each new corporate environment can have an isolated fabric and computing environment including IP space and AD Forests. Hyper-V provides services for Linux and other non-Microsoft infrastructures.

Consolidating IT infrastructures into the corporate fold allows the use of common IT resources and management while reducing redundancies and services overhead.

Multi-Tenancy provides a mechanism to merge IT resources into a corporate infrastructure in a simpler and less expensive manner. These foreign infrastructures can be merged into a common hardware computing platform, while the software infrastructure can be consolidated over time.

5.8 Test Environments

Virtualization, Multi-tenancy and Hyper-V Virtual Networking provide an environment which greatly simplifies creating a duplicate test environment. This environment can include an exact copy of the product environment including the network address space, the Active Directory domain configurations and any or all of the running workloads.

This environment can be used to test new systems, updates or software application upgrades.

Once an environment is tested, the new environment can be snapshotted and upgraded, or the user base can be swung over to from one tenant environment to another or rolled back if necessary.

6 Workloads

6.1 IaaS Service Provider

The IaaS workloads described here is the Public Cloud hoster providing services to other organizations. Since the service is consumed by other organizations without any visibility of the plans for each customer, the service provider must have a management system which reacts to usage changes quickly. This system must be sophisticated and automated.

The IaaS Service Provider has many challenges

- Large Scale
- Unpredictable Usage
- Over Subscription
- Capacity Expansion

Large Scale

The public hosting company may have hundreds of customers and many thousands of VMs and Virtual Disks. This large scale provides challenges for management systems, capacity scale out, network core capacities and performance requirements.

Unpredictable Usage

Since customers may use resources at their own discretion, CPU, Storage and Network utilization will vary at any time. Technologies such as live migration of VMs and VHDs can redistribute workloads dynamically. Management systems such as System Center Operations Manager in combination with System Center Orchestrator can be used to automate this process.

Over Subscription

Since many of customers VMs are dormant much of the time, service providers will over subscribe the infrastructure. Over subscribing resources provides a major cost savings which equates to larger profits.

Managing oversubscription requires monitoring VM utilizations and reacting quickly to capacity requirements.

Capacity Expansion

A large cost of doing business for a service provider is the hardware costs required to provide services. Purchasing hardware far in advance of need decreases the service provider’s profits. New capacity must be acquired and in operation quickly.

Described here are two techniques to add compute and storage capacity.

Combined Compute and Storage Racks

Iron Networks has developed products which combine compute and storage into a single full-sized rack. The ratio of compute to storage is based on a typical usage ratio for public hosting providers.

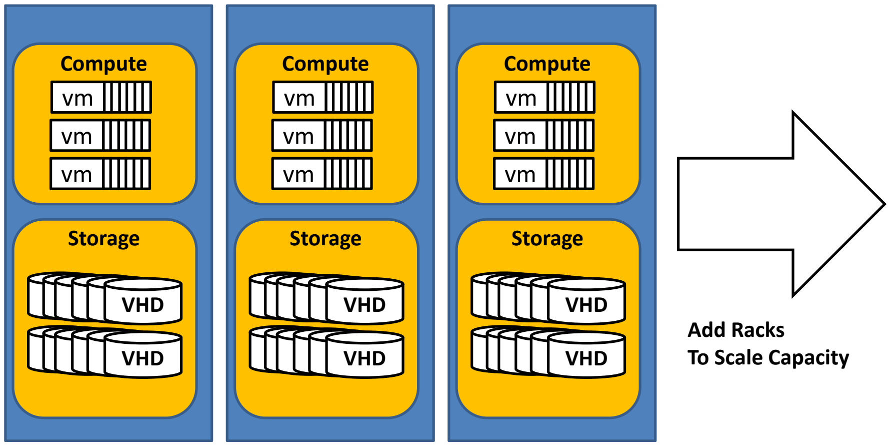

The advantages for this model include:

- Capacity is simple to increase by adding a rack at a time. Each rack contains a fully outfitted IaaS infrastructure including switching fabric, virtual networking, compute, storage and L3 aggregate routing.
- Much of the traffic between compute and storage can remain within the rack. This reduces the network load on core network infrastructures.

The disadvantages are:

- For some workloads, the ratio of compute to storage may not match the standard product as shipped. When scaling out this infrastructure and additional compute or storage is required, compute or storage racks will be required to correct the ratio.

Separate Compute and Storage Racks

Iron Networks has a family of products (model 600 series) which include full-sized racks of compute and racks of storage. This allows the service provider to add compute racks or storage racks as needed.

Compute and storage modules can be purchased separately and added to racks as needed. This allows the provider to add capacity in smaller increments.

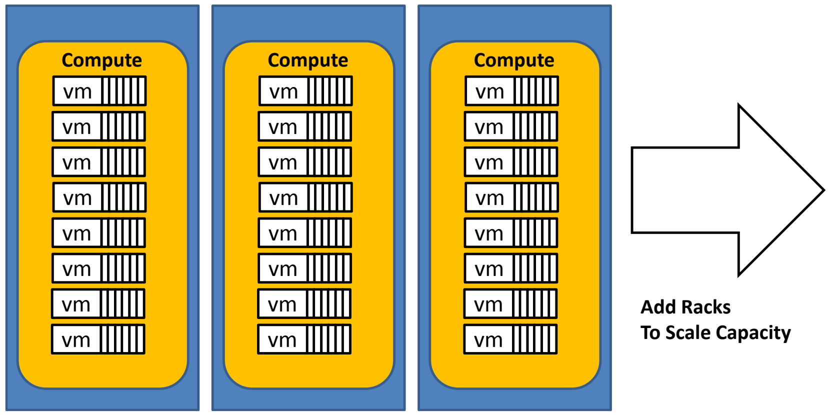

Compute racks.

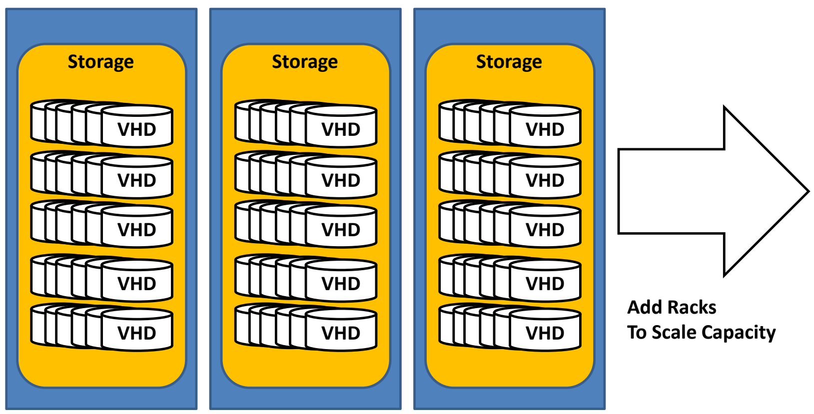

Storage Racks

Advantages are

- Capacity can be added in smaller increments
- Can add compute and storage in any ratio required

Disadvantages are

- Higher load on core network switches

6.2 VDI

Virtual desktop infrastructure (VDI) enables IT staff to deploy desktops in virtual machines on centralized, data center hardware. A centralized, optimized virtual desktop enables users to access and run their desktop and applications wherever they may be. By using virtual desktops, the IT department is able to build a more agile and efficient IT infrastructure. Flexible desktop scenarios that are running Windows operating systems give organizations the ability to choose the client computing scenarios that meet the unique needs of their businesses.

6.2.1 Persistent Virtual Desktops

A persistent virtual desktop is one in which configurations and personalization persists after each session. Persistent desktops user's profile and documents are on a separate disk specific for that user. On persistent VDI, desktop shortcuts and other user specific settings remain. Additional software packages may be installed, specific for the user.

Advantages:

The persistent VDI method operates with the same functionality as a personal PC.

Disadvantages:

- Persistent desktops have a higher administrative and support overhead.
- The amount of storage required is higher.
- Patching and upgrades are now necessary for each persistent image

6.2.2 Non-Persistent Virtual Desktops

Each virtual machine is assigned from a pool of resources when a user logs in. That virtual machine will use a user's roaming profile in order to establish their personalized settings. When you log into a virtual desktop and do your work for the day, when you log off, changes disappear from session to session. At the end of each session, the virtual desktop reverts back to its original state before login.

Advantages:

- Patches and upgrades are more easily applied to a single image in a resource pool of virtual desktops which cuts the man hours required for management.
- Centralized management of desktop images
- Less support overhead

Disadvantages:

- Sync of roaming profiles between multiple sessions is required
- Increased login time and overhead

6.3 OS365 – Exchange, Lync, SharePoint

Iron Networks have customized workload product options, designed to provide a large scalable private "Office 365-like" private IaaS infrastructure. This hardware infrastructure is developed and integrated specifically for the core Office 365 components Lync, SharePoint and Exchange. This private cloud IaaS provides the workload specific balance of CPU, Memory, Storage and Performance characteristics for each core service.

Each of these services has different requirements:

- Exchange is storage volume intensive with low storage IOPS per disk requirements. The Exchange storage farm must be constructed specifically for exchange. Exchange provides data resiliency and replication at the application level so data parity and mirroring at the infrastructure level is unnecessary.
- Lync is compute and storage performance (IOPS) intensive with less storage volume required. The storage system requires more SSD storage to meet the performance requirements.
- SharePoint is largely SQL Server based and requires a mid-level of storage volume and performance.

Iron Networks has pre-defined and integrated solutions ranging from 5000 users to 150,000 and above.

7 Solution Overview

7.1 Iron Networks IronPOD Series

IronPOD is a family of Performance Optimized Datacenters (POD) products. Every hardware and software component has been integrated and tested to provide a robust, high-performance and low cost self-contained infrastructure for VM hosting.

The product line ranges from single chassis systems for office and branch deployments to multiple rack systems.

The X100 system is a two node clustered system based on Cluster-in-a-Box (CiB) technology.

The IronPOD 600 series is designed to scale-out with dedicated compute racks and storage racks. Capacity of compute or storage can be increased by adding additional product racks.

The IronPOD 400 and 500 series products contain the compute and storage as well as the entire fabric and fabric management infrastructure to provide VM hosting.

- Core Premises Network: Infrastructure provider core network
- Aggregate Network: Switches and Virtual Networking Gateway to provide L2 services as L3 routing, failover for ingress, Site-to-Site VPN routing to remote locations.
- Tenant Network: Multi-Tenancy Virtual Networking infrastructure and tenant access to the compute nodes
- Compute Nodes: 4 to 32 node options with multiple CPU, Memory and storage options
- Storage Fabric: Storage switches providing SMB3 protocols, RDMA, RoCE and L3 RRoCE services

7.2 Iron Networks Windows Storage System Blocks (WSSB)

Iron Networks manufactures storage systems integrated with Microsoft Windows Storage Server 2012. Every hardware component is selected, integrated and tested for optimized storage performance and compatibility with Windows Storage Server and Windows Storage Spaces advanced feature sets.

There are two families of storage products, the WSS family and the ESA Family

- IRON WSS Series storage head is a custom-built, self-contained hardware appliance with dual node Cluster-in-a-Box (CiB/HA) configuration for maximum performance and reliability. Supports Storage Spaces and SMB Direct RoCE w/ 56GbE multi-channel Ethernet, and offers the fastest storage solution in the market.
- IRON ESA Series Storage is designed with two, three and four storage controller head configurations and two and four JBOD storage expansion shelves. Each support up to 60 SAS disks. The shelves can be mirrored to provide redundant enclosure high availability.

The product families range in size from single chassis products effectively priced for branch office and small business environments to multi-rack solutions.

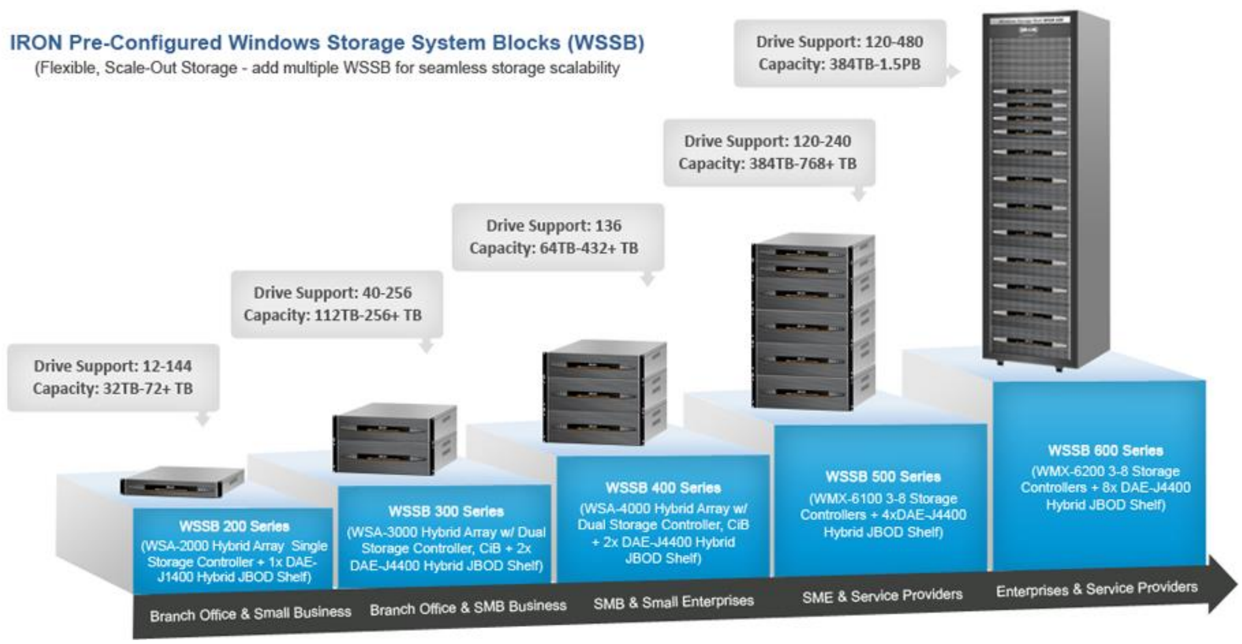

The following diagram describes an ESA storage system with four storage controller heads and four JBOD enclosures. Other configurations exist including

- Configurations with two and three controller heads
- Configurations with fewer or more JBOD enclosures
- Configurations fully populated with SSD drives for high IOPS performance

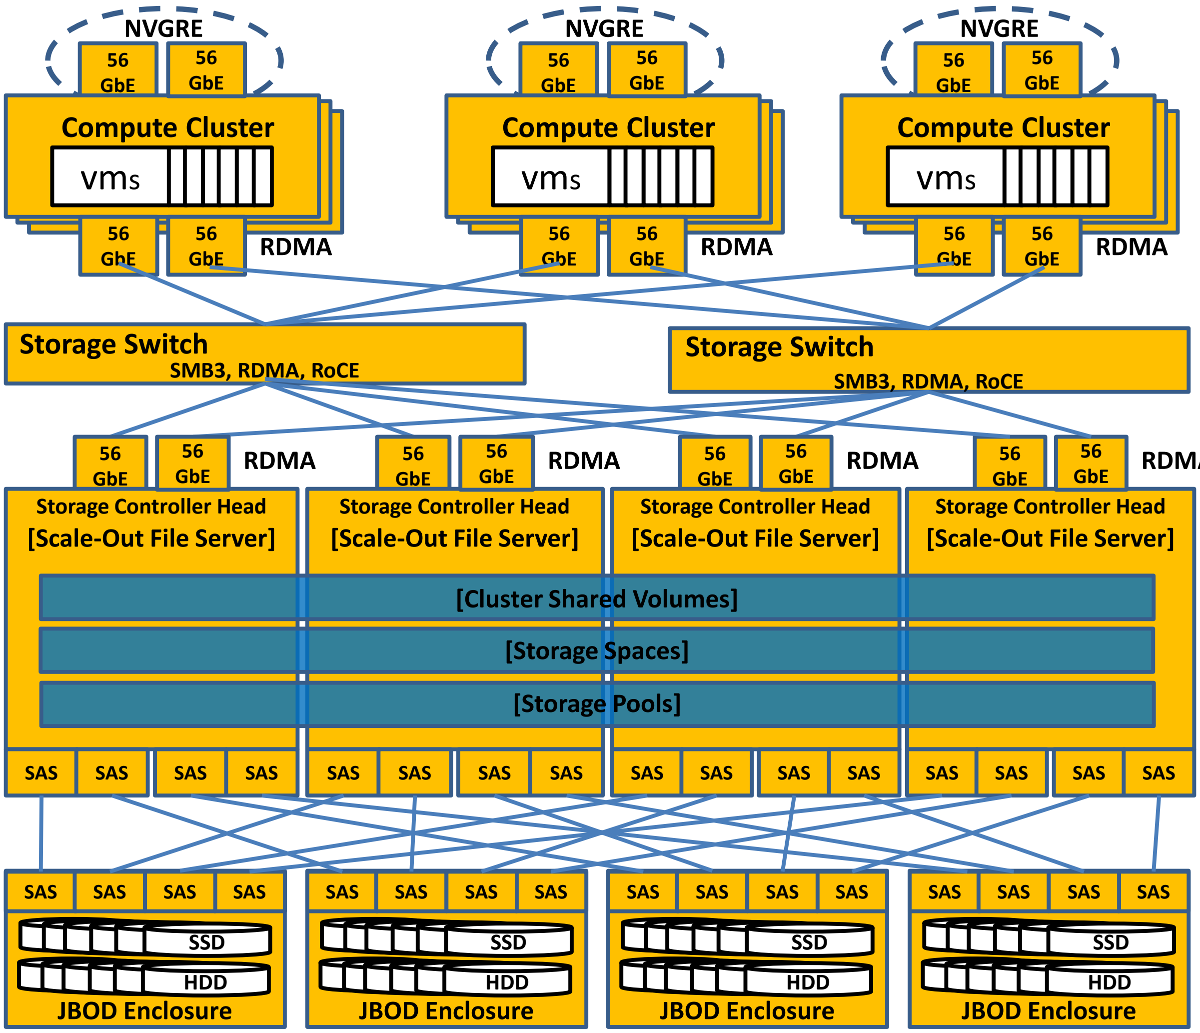

This figure describes a logical representation of a storage product with four storage controller heads (SCH) and four JBOD storage expansion shelves. Each storage expansion shelf contains and SSD disk tier and an HDD tier. Many other configurations are available.

7.2.1 Storage Availability through MPIO

The Microsoft MPIO framework allows Windows to manage and efficiently use up to 32 paths from storage devices and applications/operating systems. Although both MPIO and WSFC result in high availability and improved performance, they are not equivalent concepts. While clustering provides high application availability and tolerance of server failure, MPIO provides fault tolerant connectivity to storage. By employing MPIO and WSFC together as complimentary technologies users are able to mitigate the risk of a system outage at both the hardware and application levels.

MPIO provides the logical facility for routing I/O over redundant hardware paths connecting server to storage. These redundant hardware paths are made up of components such as the cabling, Host Bus Adapters (HBA's), Switches, and Storage Controllers and possibly even power. MPIO solutions logically manage these redundant connections so that I/O requests can be rerouted in the event that a component along one path fails.

As more and more data is consolidated on Storage Area Networks (SAN's), the potential loss of access to storage resources is unacceptable. To mitigate this risk, high availability solutions, like MPIO, have now become a requirement.

7.3 Cluster-in-a-Box (CiB)

The Iron Networks CiB product framework is a self-contained two node Windows host cluster with a JBOD-like storage enclosure designed into a single chassis.

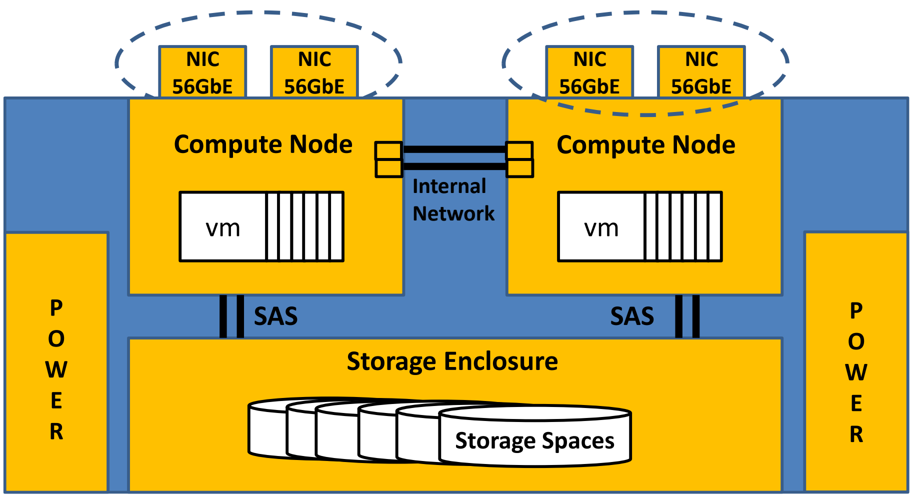

CiB Logical Diagram

Every active component in the CiB technology is redundant and hot swappable, including disk drives, compute nodes and power supplies.   Using the Windows host clustering technology, removing a compute node will dehydrate the node and all live VMs will migrate automatically to the remaining node. Once the removed node is replaced the VMs will return automatically.

The CiB product is designed into the IronPOD infrastructure for the Network Virtualization gateway (MNV), for the System Management servers (MSC) to host the System Center product VMs and as the base for several of the Iron Networks storage products (WSA).

The two node Windows 2012 host cluster configuration is designed to support Hyper-V virtual machine management as well as Windows Storage Spaces storage server.

7.4 Top of Rack (TOR) Switches

IronPOD Series 400 and above products include Mellanox SX1036 40/56GbE switches for the Aggregate, Tenant and Storage network tiers.

Mellanox network switch systems provide the highest-performing fabric solution by delivering up to 2.88Tb/s of non-blocking throughput to Enterprise data centers, with ultralow latency. It offers converged networking architecture that reduces data center costs by using a common low latency infrastructure for compute, storage and top of rack switches.

- 40/56GbE uplink for top-of-rack access networking, dual switches for HA.
- 40/56GbE for tenant network and storage network switches, dual switch for HA.
- 40/56GbE for storage networking, dual switches for HA.
- 40/56GbE dual channel (failover) for network adapters.

Specifications and Performance		
Model Number	SX1036	
Ports	36	
Backplane Speed	2.88 Tbps	
L2 Latency	223 nanoseconds	(all frame sizes)
L3 Latency	333 namoseconds	(all packet sizes)

7.5 Microsoft System Center (MSC)

Fabric Management for the IronPOD product infrastructure is provided by the MSC product family. The MSC series management appliance is a self-contained hardware platform designed to run independently of the IronPOD compute and storage tiers. This independence and isolation from the workload tiers is important to provide:

- Bare-Metal provisioning to all components in each rack. This requires providing DHCP services and other services, and from the management system, administrators can bring individual components or entire racks to a factory defaults configuration.
- Active Directory services for internal rack management.
- The System Center family of products runs on the MSC modules. This allows System Center to run independently of the customer workload in its own security domain and not complete for resources.

The Microsoft System Center 2012 Suite offers an integrated management platform for easy and efficient management of datacenters and hybrid cloud environments. It serves as a single platform for comprehensive management of applications, services, physical resources, hypervisors, software defined networks, configuration, and automation.

Microsoft® System Center™

Iron Networks MSC Series management head is a custom-built turnkey appliance with dual node Cluster-in-a-Box (CiB) configuration for maximum performance & reliability.

When the IronPOD product racks are deployed, the Network Virtualization Gateway and every compute and storage component is “raw” and unconfigured. The only module operational is the MSC Fabric Management system.

On initial deployment of the MSC system, the MSC boots a factory default image and requests several preliminary parameters such as Active Directory domain details and IP subnet addresses. The MSC then builds itself into a fully operational management cluster including the Deployment system, the Active Directory servers and the System Center suite of Fabric Management services.

Management Node 1 CIB Node 1	Management Node 2 CIB Node 2
Active Directory Server 1	Active Directory Server 2
WDS Server (PXE Boot)	Deployment Server
SC VMM Node 1	SC VMM Node 2
SQL Cluster Node 1	SQL Cluster Node 2
Configuration Manager	SC SPS Web Server
SC Operations Manager Node 1	SC Operations Manager Node 2
SCOM DB SQL Server 1	SCOM DB SQL Server 2
Orchestrator Server	SSRS (SCOM Reporting)
Orchestrator Row Book Server 1	Orchestrator Row Book Server 2
SC Control Node	
Firmware Manager	

The figure above describes a logical representation of the MSC system and the services running on each node.

7.6 Microsoft Network Virtualization (MNV)

Microsoft Network Virtualization Platform Solution Overview

Hyper-V Network Virtualization creates "Virtual Networks" similar to how server virtualization (hypervisor) creates "Virtual Machines". Network virtualization decouples and isolates virtual networks from the physical network infrastructure and removes the constraints of VLAN and provider IP address assignment from virtual machine provisioning.

This flexibility makes it easy for customers to move workloads to IaaS clouds and efficient for hosters and datacenter administrators to manage their infrastructure, while maintaining the necessary multi-tenant isolation, security requirements, and supports overlapping IP address ranges.

Enables flexible workload placement – Network isolation and IP address re-use without VLANs

Hyper-V Network Virtualization decouples the customer's virtual networks from the physical network infrastructure of the hosters, providing freedom for workload placements inside the datacenters. Virtual machine workload placement is no longer limited by the IP address assignment or VLAN isolation requirements of the physical network because it is enforced within Hyper-V hosts based on software-defined, multitenant virtualization policies.

Virtual machines from different customers with overlapping IP addresses can now be deployed on the same host server without requiring cumbersome VLAN configuration or violating the IP address hierarchy. This can streamline the migration of customer workloads into shared IaaS hosting providers, allowing customers to move those workloads without modification, which includes leaving the virtual machine IP addresses unchanged. For the hosting provider, supporting numerous customers who want to extend their existing network address space to the shared IaaS datacenter is a complex exercise of configuring and maintaining isolated VLANs for each customer to ensure the coexistence of potentially overlapping address spaces. With Hyper-V Network Virtualization, supporting overlapping addresses is made easier and requires less network reconfiguration by the hosting provider.

In addition, physical infrastructure maintenance and upgrades can be done without causing a down time of customer workloads. With Hyper-V Network Virtualization, virtual machines on a specific host, rack, subnet, VLAN, or entire cluster can be migrated without requiring renumbering or major reconfiguration.

Enables easier moves for workloads to a shared IaaS cloud

With Hyper-V Network Virtualization, IP addresses and virtual machine configurations remain unchanged. This enables IT organizations to more easily move workloads from their datacenters to a shared IaaS hosting provider with minimal reconfiguration of the workload or their infrastructure tools and policies. In cases where there is connectivity between two datacenters, IT administrators can continue to use their tools without reconfiguring them.

Enables live migration across subnets

Live migration of virtual machine workloads traditionally has been limited to the same IP subnet or VLAN because crossing subnets required the virtual machine's guest operating system to change its IP address. This address change breaks existing communication and disrupts the services running on the virtual machine. With Hyper-V Network Virtualization, workloads can be live migrated from servers running Windows Server 2012 in one subnet to servers running Windows Server 2012 in a different subnet without changing the workload IP addresses. Hyper-V Network Virtualization ensures that virtual machine location changes due to live migration are updated and synchronized among hosts that have ongoing communication with the migrated virtual machine.

Enables easier management of decoupled server and network administration

Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, and network administrators can focus on overall network infrastructure and traffic management. This enables datacenter server administrators to deploy and migrate virtual machines without renumbering the virtual machines. There is reduced overhead because Hyper-V Network Virtualization allows virtual machine placement to occur independently of network topology, reducing the need for network administrators to be involved with placements that might change the isolation boundaries.

Simplifies the network and improves server/network resource utilization

The rigidity of VLANs and the dependency of virtual machine placement on a physical network infrastructure results in overprovisioning and underutilization. By breaking the dependency, the increased flexibility of virtual machine workload placement can simplify the network management and improve server and network resource utilization. Note that Hyper-V Network Virtualization supports VLANs in the context of the physical datacenter. For example, a datacenter may want all Hyper-V Network Virtualization traffic to be on a specific VLAN.

Compatible with existing infrastructure and emerging technology

Hyper-V Network Virtualization can be deployed in today's datacenter, yet it is compatible with emerging datacenter "flat network" technologies.

Provides for interoperability and ecosystem readiness

Hyper-V Network Virtualization supports multiple configurations for communication with existing resources, such as cross premise connectivity, storage area network (SAN), non-virtualized resource access, and so on. Microsoft is committed to working with ecosystem partners to support and enhance the experience of Hyper-V Network Virtualization in terms of performance, scalability, and manageability.

7.7 Microsoft Network Virtualization Gateway

The MNV Gateway resides between the Tenant network switches and the Aggregate switches and in an integral part of the Tenant network fabric.

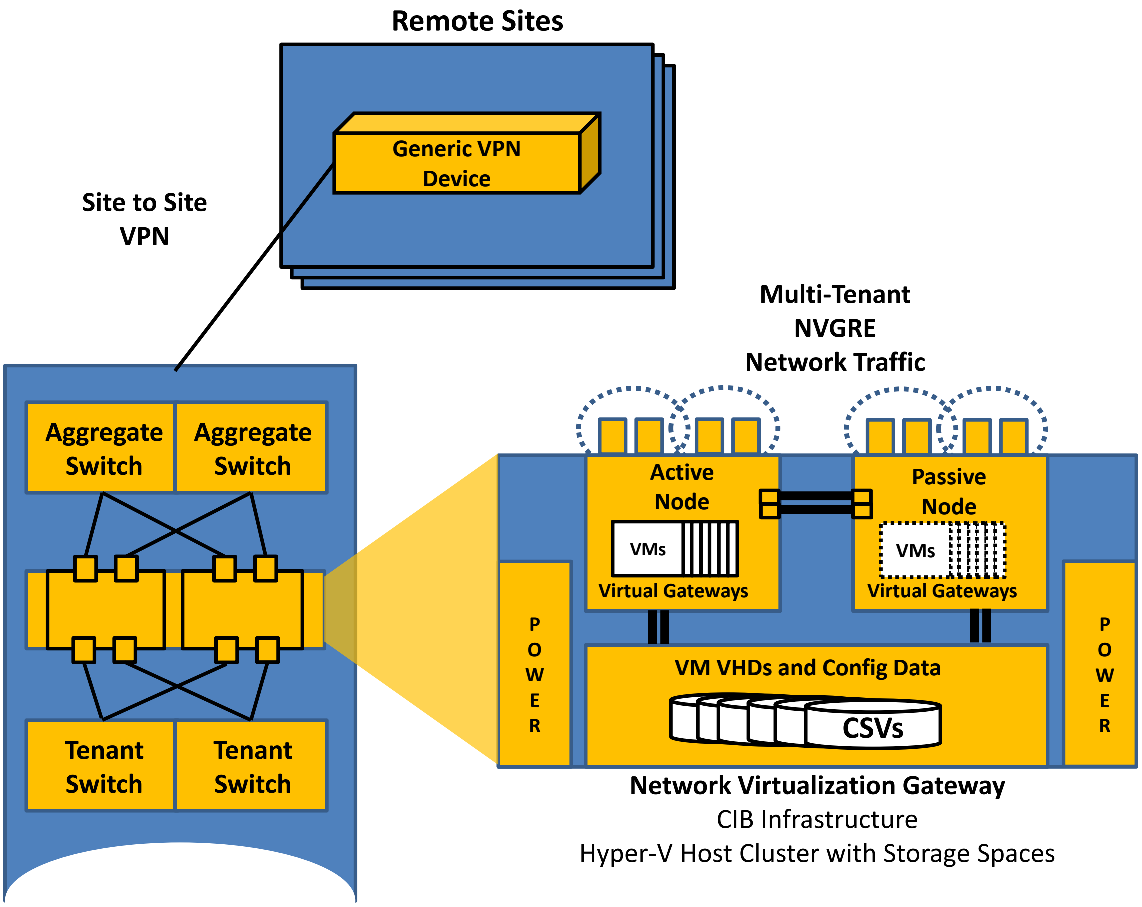

Key Features

Site to Site VPN: The MNV Gateway connects Site to Site VPN services between the Tenant VMs running inside Tenant Virtual Networks and remote tenant site locations.

NVGRE Offload Hardware: Special hardware is integrated into all compute nodes and the MNV Gateway to provide hardware offloads. NVGRE traffic is processed inside the hardware to increase network performance, allow LBFO failover and provide scalability.

Active/Passive Cluster: The Virtual Network gateway hardware is based on the Cluster-in-a-Box cluster technology where a complete Windows host cluster is integrated into a single hardware product. The cluster operates in Active/Passive cluster mode.

Virtual Gateway VMs: Inside the MNV Gateway are multiple VMs. The VMs provide virtual gateway services for tenant networks.

Multi-Tenant NVGRE Network Traffic: NVGRE is the IP encapsulation protocol which provides the isolation between tenant networks.

Hyper-V Host Cluster and Storage Spaces: The infrastructure supporting the gateway functions is based on Hyper-V Clusters and contains a full Storage Spaces implementation to support VM mobility and centralized configuration services.

7.8 Compute Tier

The compute tier consists of a group of four-plex chassis with compute node servers. The compute nodes are Hyper-V enabled and creates a fabric which scales in the form of a large-scale failover cluster.

Windows Server 2012 supports up to 64 nodes in a single cluster.

Every component in the compute nodes are hot swappable including compute blades, power supplies and hard disks. In the case of server blade failures, Windows clustering will dehydrate the node automatically. Iron Networks provides several service offerings from managed service monitoring and call-home where a replacement server component will be automatically shipped to advanced replacement services.

Specifications		
CPU Core	6 to 16	
Memory	32-512GB	
GPU	Optional	VDI Workloads

8 References

Microsoft Private Cloud Fast Track Reference Architecture
http://download.microsoft.com/download/3/4/D/34D3CCF9-DF32-48BD-AD99-A004433D86FE/Microsoft%20Private%20Cloud%20Fast%20Track%20Reference%20Architecture%20Guide.pdf

Microsoft Private Cloud Fast Track Reference Deployment Guide
http://www.microsoft.com/en-us/download/details.aspx?id=30422

Microsoft Private Cloud Fast Track Reference Operations Guide
http://www.microsoft.com/en-us/download/confirmation.aspx?id=30416

Final Version of NIST Cloud Computing Definition Published
http://www.nist.gov/itl/csd/cloud-102511.cfm

NIST Definition of Cloud Computing
http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Information Technical Instruction Library
http://www.itil-officialsite.com/

Microsoft Private Cloud Reference Model
http://social.technet.microsoft.com/wiki/contents/articles/4399.private-cloud-reference-model.aspx

Hyper-V Dynamic Memory Configuration Guide
http://technet.microsoft.com/en-us/library/ff817651(WS.10).aspx

Windows Server High Availability with Microsoft MPIO
http://www.microsoft.com/en-us/download/confirmation.aspx?id=9787

Hyper-V Live Migration Network Configuration Guide
http://technet.microsoft.com/en-us/library/ff428137(WS.10).aspx

Microsoft Operations Framework 4.0
http://technet.microsoft.com/en-us/library/cc506049.aspx