bigstep®

The Full Metal Cloud

The world's highest performance cloud purpose-built for big data.

bigstep.com

FULL METAL CLOUD

The Full Metal Cloud is the world's highest performance cloud, purpose-built for big data. It is a bare-metal IaaS aimed at high performance applications such as big data analytics, high volume batch processing, real time analytics, high volume stream data analysis, machine learning and natural language processing.

PERFORMANCE LIKE NO OTHER

Time it takes to read/write 1 TB in memory - smaller is better

THE FULL METAL CLOUD HAS THREE ESSENTIAL COMPONENTS:

- **Bare metal infrastructure**
  The actual metal (servers, storages, and networking) that applications are deployed on. We currently use HP servers as compute instances.

- **Big Data on Full Metal application ecosystem**
  Powerful big data applications are integrated with the orchestration layer and have fully automated deployment, configuration, and scaling processes.
  The result is a comprehensive application ecosystem which can be accessed at the click of a button and answers most big data use cases.

- **Orchestration layer & Control Center**
  The brain behind the Full Metal Cloud, the orchestration layer cloudifies both the metal, as well as the applications running on top of it (Hadoop, NoSQL, search & analytics engines) and is fully controllable via the API or Control Center. The frontend of the orchestration layer is the Control Center, the Full Metal Cloud's management UI, which enables users to deploy complex big data setup with simple, drag & drop actions.

Datameer · hadoop · elastic · EXASOL · DATASTAX · cloudera · Couchbase · splunk>

Got a question? Look no further. Call us now at: **+44 (0) 207 510 9298** or say **hello@bigstep.com**

FULL METAL COMPUTE INSTANCES - FMCI

There are two types of instances: Full Metal Compute Instances and Storage Instances, which host Full Metal Drives.

- **Full Metal Compute Instances**
  FMCIs are diskless physical servers highly optimized for performance. FMCIs have varying memory and CPU configurations: from 8 to 192 GB of RAM, and from one quad-core CPU to two ten-core CPUs per physical instance. We can also provide compute instances with local disks but recommend that they be used only as scrub disks.

  Each instance has at least 5 physical Ethernet ports. Four are network ports and one is for Out Of Band / OOB (IPMI and KVM over IP) connections. Network ports are controlled by the client and can be aggregated, as well as allocated to SAN, LAN, or WAN traffic. Single CPU configurations have 4 ports of 1 Gbps + OOB, while dual CPU configurations have 4 ports of 10 Gbps + OOB.

- **Storage Instances**
  They are specifically designed for high I/O and throughput, and have either all-SSDs or all-disk drives. Each storage instance has a minimum of 40 Gbps connectivity in full duplex and delivers over 1 million IOPS. Compute instances boot from drives located on storage instances.

CLOUDIFYING METAL

Users can order and manage compute instances and drives. The orchestration layer allocates drives to compute instances from the nearest storage instance, so that traffic is contained, as much as possible, within the same rack.
Compute instances are managed in groups, called Instance Arrays, in order to simplify management of large scale clusters for distributed big data applications.

An Instance Array can contain a single compute instance or as many as needed. Changes applied to an Instance Array apply to all compute instances it contains, at the same time. This includes changes in hardware specs, network, storage, or deployed software. Compute instances can be managed separately using SSH/RDP or the OOB interface.
Deploying an Instance Array with the operating system installed takes about 2 minutes. Drives are also managed in groups, called Drive Arrays.

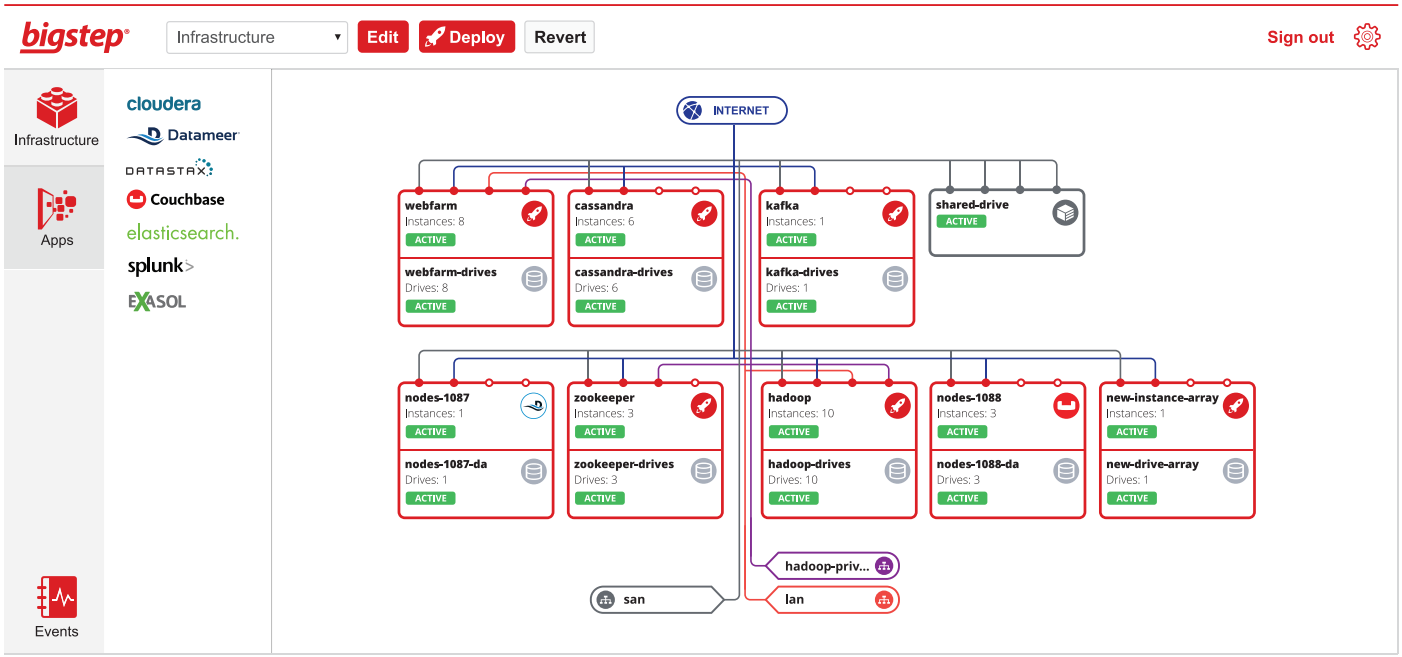

FULL METAL STORAGE

Storage throughput is paramount to performance in big data. As our all-SSD storage instances can yield at least a million IOPS and 5GB/s throughput, the placement of the storage in close proximity to compute instances becomes critical. We use state of the art wire-speed guaranteed 10 GbpE cut-through switches which deliver nanosecond latency between compute instances and the storage.

We also use Ethernet cards, capable of RDMA and offloading TCP and iSCSI traffic, in order to achieve even lower latencies and higher throughput. Our system allocates storage drives in close proximity to compute instances (usually within the same rack), to keep latency at a minimum.

Compute instances have either 4 x 1 Gbps or 4 x 10 Gbps connections. Ports can be aggregated - the maximum storage throughput achieved by a single compute instance is 40 Gbps full duplex. On compute instances with local drives, these should be used only as scrub disks, as data on them would not survive an instance upgrade/downgrade or migration. Their content is zero-filled when the physical server is returned to the available pool. For this reason, we highly recommend the use of distributed block storage or object storage.

FULL METAL NETWORK

The Full Metal Network is one of the key differentiators of the Full Metal Cloud. It is a fully automated software defined network. Switch ports are configured instantly based on users' API calls or actions in the Control Center.

When an API call is made or an action is performed in the Control Center, the Full Metal Network automatically creates a secure tunnel between the two physical ports, linking them via an isolated layer 2 pseudo-wire connection. This tunnel is a bare metal environment in which VLANS can be driven and traffic can be broadcast. When a new port is added to the same LAN, the layer 2 network is expanded, just like a physical on premise network.

The Full Metal Cloud uses a flat network architecture, with almost no aggregation layer and it is optimized for east-west traffic as well as north-south. Each switch on the path between instances is guaranteed to sustain line rate performance on all links at all times.

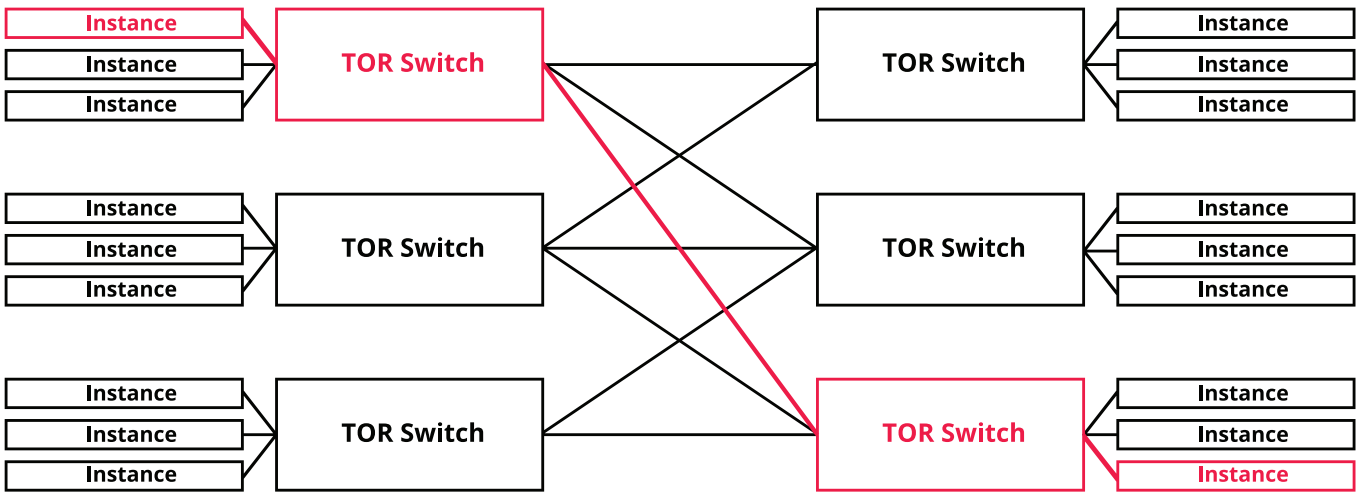

Designing a complex network architecture in the Control Center is as simple as drag & drop. The interface allows building networks as easy as drawing a wireframe on paper. It takes literally minutes to:

- Create the physical LANs
- Define WAN connections
- Cross-connect Instance Arrays
- Aggregate ports
- Define secure zones

FULL METAL IP PATHFINDER

IP Pathfinder is the automated IP management service which ensures that scaling vertically, horizontally, and instance migrations work seamlessly:

- Persistent public IPv4 subnets are automatically allocated for each Full Metal Compute Instance during the boot process.
- IPv6 addresses are allocated per each Full Metal Compute Instance – making every infrastructure in the Full Metal Cloud IPv6 ready
- IPv4 and IPv6 addresses allocated to Full Metal Compute Instances are transported to any physical server or drive allocated to the instance, during scaling or migrations

OPERATING SYSTEMS

Operating systems in the Full Metal Cloud need to read the iSCSi targets and network configuration from the iBFT table. We support CentOS, Ubuntu, Windows, VMWare ESXi, REHLs and many other OSs. We also provide deployment templates for CentOS and Ubuntu, while Windows will follow shortly.

SECURITY ON FULL METAL

DATA SECURITY

The Full Metal Cloud provides isolated environments for hosting critical data, with the following characteristics:

- **Bare metal, isolated, single-tenant servers**
  You don't share hardware with anyone; security risks associated with multi-tenant virtualized environments are completely eliminated.

- **Root access & data security**
  After the initial provisioning, root access credentials can be changed, preventing anyone, including Bigstep technical staff, from accessing your systems.
  Data on the disk can be encrypted. Existing security procedures such as Active Directory integration or periodic patches can be applied to the instances.

- **Isolated layer 2 broadcast domains**
  Each client network is completely isolated at the physical level. Switch ports are physically dedicated to individual client accounts.
  Traffic cannot be intercepted or injected in any way into client controlled networks.

- **Dedicated physical network**
  You are in complete control of the network, at the physical level. There is no virtual switching involved and no multi-tenancy.

- **WAN traffic isolation**
  Public connections work on isolated layer 2 broadcast domains up to the gateway, so no interception or unauthorized traffic injection is possible.

- **LAN networks**
  They are physically isolated, so users can configure their own private IPs and drive all their VLANs.
  There cannot be any private IP or VLAN conflict with another user in the Full Metal Network.

- **The Storage Network**
  This is a key security area. SAN traffic is isolated in multiple ways:
  - Traffic to Compute Instances is routed only inside the subnet allocated to each instance, preventing access to/from any other device in the Full Metal Network. Corresponding switch ports are used as layer 3 gateways for the SAN traffic, to prevent traffic sniffing.
  - Layer 3 ACLs are in place to filter traffic and allow only specific storage communication to flow from/to specific instances in the network.
  - PXE DHCP messages used to authenticate an instance are based on a tag the switch adds to all DHCP traffic coming from a port. This ensures that there is no way for a user to impersonate another.

- **WAN traffic**
  WAN traffic sent via public IPs between compute instances is transferred inside a layer 2 broadcast domain. Each user has a subnet and layer 3 gateway interface in the layer 2 broadcast domain, so a user's WAN traffic can never be intercepted by others.

- **RADIUS authentication**
  Operations our orchestration layer performs on switches are regulated by RADIUS authetication.

- **Comms encryption**
  API and orchestration layer communications are encrypted, signed with timestamps, and secured against reply attacks, identity theft (impersonation) or request forgery.

- **Private data encryption**
  All critical data is encrypted. Passwords are not stored in the database.

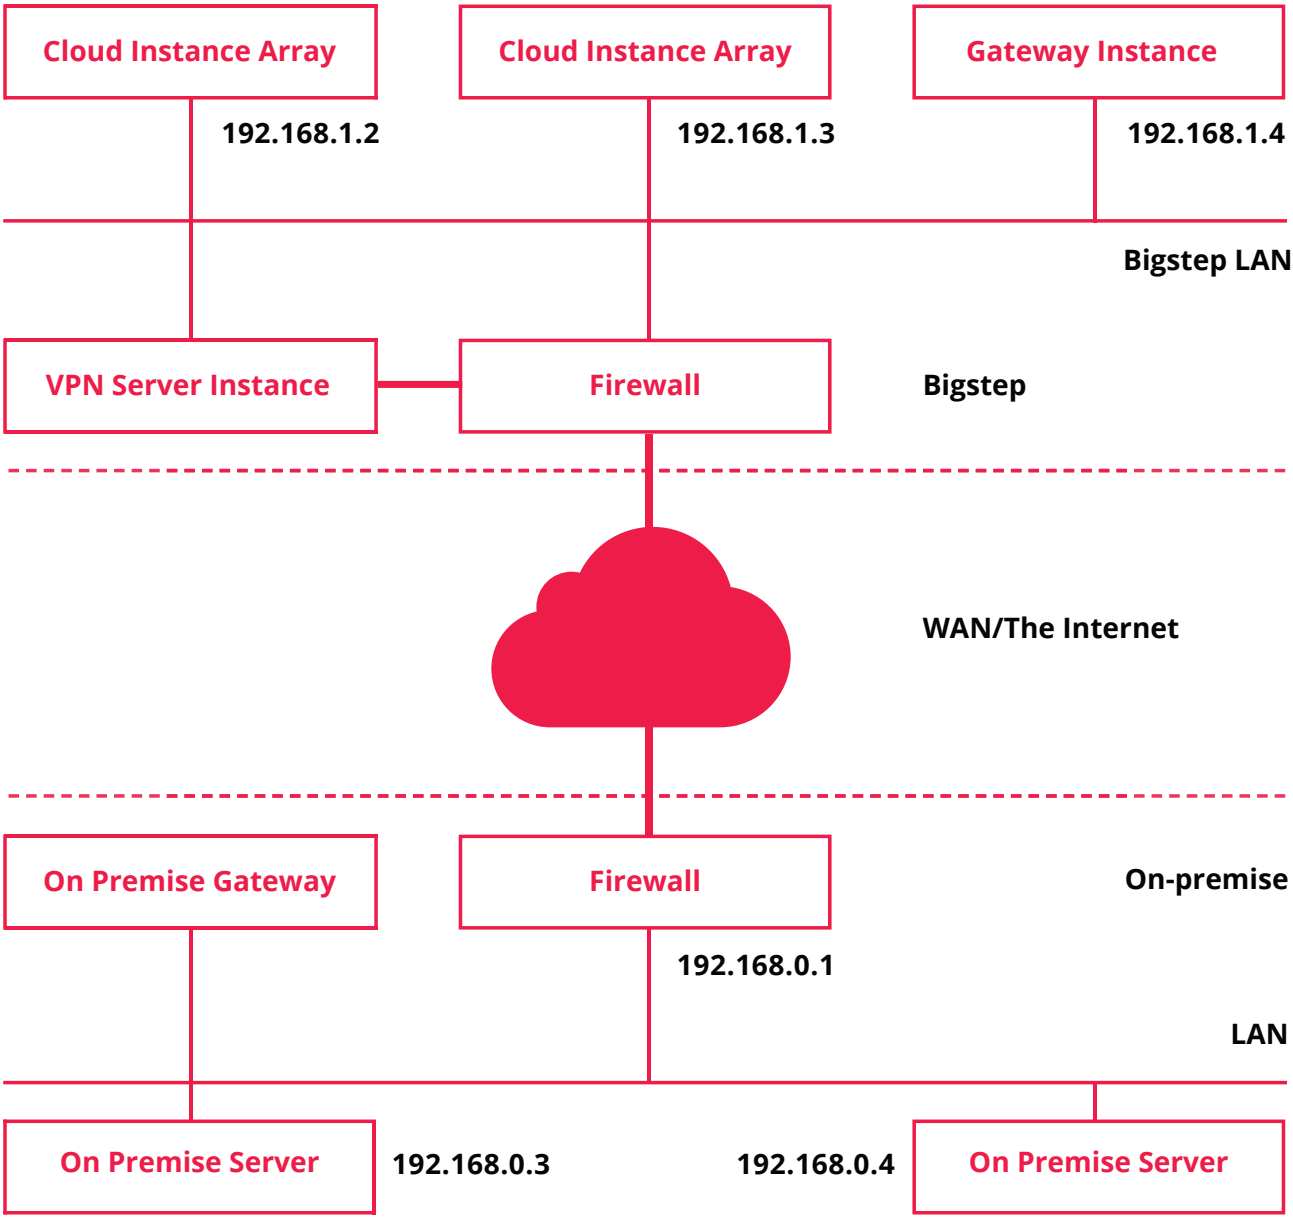

DC & PHYSICAL SECURITY

The data centers are Tier-3 certified. Sites are manned 24/7 by dedicated security teams. Physical security barriers around the perimeter of the site, 24/7 monitored CCTV and access card restricted multi-layer access points also guarantee that data is kept safe at all times.

- 24/7 on-site dedicated security teams
- 24/7 live monitored CCTV
- Multi-layer access control system with man trap access restriction

CERTIFICATIONS & COMPLIANCE

- Bigstep is ISO 27001 / 2013 certified and in the process of being PCI DSS certified.
- The Full Metal Cloud is compliant to the new GDPR and the NIS Directive.
- We work with our clients in achieving certifications or internal audits where required.

CLOUD VERSUS ON-PREMISE TCO

The idea of the cost effectiveness of running on-premise hardware is still very much alive despite evidence to the contrary. The primary reasons for that are the hidden costs of on premise, which are either ignored or too complex to calculate.

MOORE'S LAW ADJUSTED TCO

History has shown us that advances in technology modelled by Moore's law, as well as competitive pressure, forces all the cloud providers to lower pricing with roughly 30% per year. This yields a TCO reduction at the end of the regular 5 year depreciation period of roughly 40% over regular linear depreciation.

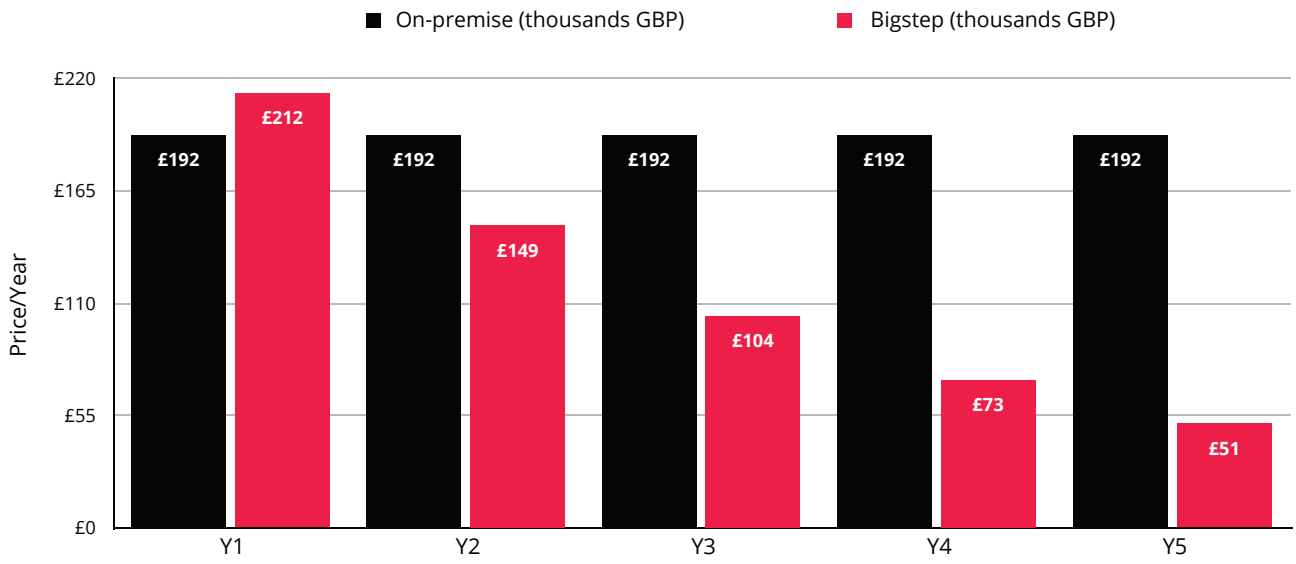

This simulation is based on the cost of owning a rack of 30 servers in either on-premise or on Bigstep.

SIMILAR EQUIPMENT, DIFFERENT USAGE PATTERNS

The cloud providers and thea average enterprise buy roughly the same equipments but what differentiates them is:

1. Economies of scale
2. Optimal use of network resources, connectivity, cooling and datacenter space
3. Optimal employment of personnel for datacenter operations, logistics and acquisitions
4. High degree of automation
5. High consolidation and usage ratios across all resources

HIDDEN COSTS OF ON-PREMISE

When calculating TCO for on-premise don't just add the price of the servers, but also consider:

- Networking costs: TOR switches, aggregation switches, optics, upgrades to the backbone if needed
- Power consumption: i.e a server uses 400-1000W of power
- Cooling costs: an extra 40% of the power is used for cooling
- Setup costs & next business day services
- Orchestration software costs - including consultancy, development, systems integration
- Datacenter floor cost
- Datacenter upgrade costs: generators, cooling equipment, fire suppression systems, etc.
- 24/7 NOC related costs
- Personnel: datacenter maintenance, logistics, overtime payments due to weekend or nighttime works
- Financing costs

THE TIME RESOURCE

Time is by far the most ignored of all the hidden costs. If a project takes three months to implement instead of half an hour it costs the business:

- Three months worth of salaries for everybody involved
- Three months worth lost potential income generated by the infrastructure
- Lost income due to late arrival of a product based on that infrastructure

THE IMPLICATIONS OF FLEXIBILITY

Bringing in a new server in an enterprise involves:

- Budged approval
- Procurement approval
- Shipment
- Installation
- Changes to the network
- Management software updates

The cloud changes all that with its key feature: flexibility. Scalability plateaus can now be adjusted every day during those 5 years, making guestimation a thing of the past. This changes the way IT operates across the board.