

WHITE PAPER

Making the Business Case For Software-defined Bare Metal Infrastructure

A guide for business leaders on how to gain cloud-like agility with physical equipment and eliminate human error.

WHITE PAPER

Contents

Why bare metal automation?	4
Enable self service bare metal	4
Deliver infrastructure faster	4
Reduce resource waste	5
Why do we see a renewed interest bare metal?	6
New architectures do not need virtualization	6
Local storage is now increasingly important	6
100Gbps ethernet gear is now mainstream.....	6
Companies need to manage a mix of virtualization and bare metal.....	6
Why is infrastructure abstraction important?	7
Software defined bare metal vs Converged vs Hyperconverged?.....	7
Conclusion	7

Making the Business Case for Software-defined Bare Metal infrastructure

Many companies are increasingly relying on bare metal automation to speedup service delivery, eliminate human errors and delays and unlock new revenue streams.

Software defined, composable infrastructure enables companies to treat physical compute, storage and network as pools of resources that can be provisioned as needed. It's an emerging category of infrastructure that's aimed at optimizing IT resources and improving business agility.

It solves the "last mile problem of automation", by enabling modern centralized management and cloud-like consumption of physical equipment.

This problem is costing companies an estimated \$1,400,000 per year in man hours alone and an additional \$700,000 per year in underutilized resources for every 5000 servers¹.

This novel approach also comes amid renewed interest in bare metal fueled by the increased use of containers, novel application architectures, NoSQL databases, realtime processing and others. 53 percent of organizations now use Docker and 48 percent use Kubernetes².

Why bare metal automation?

While everything from the virtualization layer and above has been automated for many years, the bare metal layer: servers, switches, storage remain largely un-automated. We call this “*the last mile problem of automation*”.

Many companies are trying to solve this problem for a variety of reasons. When asked what is the primary motivation of using a bare metal automation software instead of manually managing server & network, users replied⁹:

1. Enable self-service bare metal (54%)
2. Deliver infrastructure faster (28%)
3. Reduce human errors (11%)
4. Reduce resource waste (7%)

Secondary objectives we hear frequently are:

- Infrastructure-as-code
- End-to-end visibility
- Supporting edge cloud scenarios

However, software defined bare metal infrastructure also offers other benefits. It can help with businesses build an abstraction layer on top of physical equipment allowing developers to consume resources without worrying about variations between environments, all without sacrificing performance.

It also helps IT departments drive innovation, supporting the R&D department with self-service, container optimized infrastructure as well as help the business drive costs down by as much as 6x when compared to cloud-based resources by using in-house, self-managed off the shelf bare metal infrastructure.

Enable self service bare metal

Interest in running workloads directly on bare metal has been rising steadily in recent years due to the increasing use of containers, NoSQL databases, ML/AI workloads, use of specialized hardware such as GPUs, Intel Optane memory, SmartNICs and others.

The Bare Metal Cloud market is now a \$2.1bn market projected to grow almost ten fold to \$20.2bn by 2027 (38% CAGR)³.

The percentage companies of using Kubernetes is also rising rapidly with 48% of companies already using it and another 25% planning to use it.

To support this, end-users are now looking to consume bare metal resources in the same manner as they do VMs in the cloud.

Naturally, many enterprises and service providers are looking to either build a new stand-alone bare metal cloud offering or augment an existing service by leveraging bare metal automation.

Deliver infrastructure faster

With the advent of cloud, the window of opportunity for delivering new services to users has been shrinking. Users now expect cloud-like “instant access” experience.

Without bare metal automation, the time to deliver a server to a client, even under the best of conditions, relying on automation of everything above the operating system level is around 14 hours⁵:

MAKING THE BUSINESS CASE FOR SOFTWARE DEFINED BARE METAL INFRASTRUCTURE

Action	Done by	Time
Server re-cabling	Remote hands	30m
Hand-over		3h
Switch L3 config.	Network eng.	30m
Hand-over		3h
Server RAID config.	DevOps eng.	10m
OS install	DevOps eng.	15m
Network & FW config.	DevOps eng.	20m
Hand-over		3h
Storage config.	Storage eng.	15m
Hand-over		3h
Verification & final stack deployment (with automation)	DevOps eng	1h
	Total	14h

In theory, all those actions take under 30 minutes but in practice in between these functions there are frequent delays as the work waits for the next engineer to pick it up. With an average hand-over time of 3h, the delivery time for a server would be 2 working days.

At the same time, the complexity of required infrastructures has been increasing. Clients demand entire clusters rather than individual servers. For an infrastructure composed of 5 servers interconnected in some fashion, the delivery time will be between 5 and 10 work days (1-2 weeks) depending of course on the dynamics of the team and degree of automation.

Software defined infrastructure virtually eliminates all this effort delivering complete bare metal solutions in a matter of minutes.

Reduce human errors

For most teams, the most difficult part of deploying any kind of infrastructure is provisioning and configuring the network:

1. Allocating IPs;
2. Configuring ports on switches;
3. Setting up the network configuration of the operating system;
4. Updating the firewall configuration;

The potential for human errors and their impact is much higher at this stage. A wrong IP configured on an interface can bring down other production servers. Using tools like TCP dump to try to understand why a certain traffic is not reaching a newly provisioned infrastructure is very common, adding to the delays in delivering a server.

Bare metal automation solves this problem by allocating IPs automatically, configuring VLANs automatically etc.

Reduce resource waste

According to Uptime institute, 20% of servers are under-utilized (even with virtualization), costing a company \$250,000 in power and another \$236,000 for every 5000 production servers¹⁰.

This is because without automation, bare metal equipment is hard to re-purpose thus companies simply choose to leave them running.

In addition, the lack of integrated, automated tracking of all resource allocation there is further waste:

1. Orphan storage LUNs
2. Lost or misplaced servers and components
3. Allocated but unused IPv4 space

Resource waste thus amounts to an estimated \$703.000

per year for 5000 servers¹.

In contrast, equipment managed by a Software-defined bare metal Infrastructure solution enables companies to easily re-purpose servers from tenant to tenant or project to project or simply shut them down, allowing companies to greatly improve resource utilization.

Automatic discovery and automatic life-cycle management also eliminate lost or misplaced equipment by keeping track of everything without human intervention during the entire life of an equipment.

Why do we see a renewed interest bare metal?

New architectures do not need virtualization

The renewed interest in bare metal is fueled in part by the changing resource consumption patterns:

- From traditional monolithic workloads to micro-services-based architectures;
- From operating systems to docker images;
- From master-slave database architectures to sharding & replication based distributed “NoSQL” databases;

The old one-VM-one-workload model is increasingly obsolete. Most applications are now distributed and include their own application specific scheduler that spreads work (or traffic) across hosts. Containers also remove the need for an OS per worker.

These new architectures are able to achieve consolidation ratios higher than VM-based workloads while also being faster, more scalable and more resilient.

Local storage is now increasingly important

Local SSDs access times are now at PCIe speeds or close to RAM speeds, making them ideal candidates for storing portions of NoSQL databases called "shards".

Since workloads such as NoSQL databases no longer need to move between hosts and incorporate their own data protection strategies such as replication they are able to use local storage directly while applications can also be made aware of data locality, avoiding cross-host network transport for every access.

For example, Aerospike can store data into SSD cells directly, bypassing the file-system altogether, perform de-duplication using the SSD's block address table etc.

A modern composable infrastructure solution will be able to leverage local storage and remote storage at the same time.

100Gbps ethernet is now mainstream

Costs per port for 25G/50G is now lower than that of a 10G¹¹. At these data volumes, virtual switching and encapsulation and decapsulation is no longer an option due to high CPU usage and latencies involved. Offloading to SmartNICs can help reduce this overhead but are expensive.

Using hardware ASICs instead of virtual switching with encapsulation enables a 2.7x increase in throughput and a 1.7x decrease in latency⁶.

Companies need to manage a mix of virtualization and bare metal.

While the new architectures make virtualization obsolete, many companies have and will continue to have legacy applications to run in parallel to the new systems.

Without automation, moving resources between these islands is difficult and error prone.

With automation, virtual private clouds such as VMWare vSphere® can be tenants of a lower-level bare-metal cloud which means they can grow and shrink as needed and the resources shifted to other clusters such as Kubernetes or OpenStack.

Why is infrastructure abstraction important?

Software-defined infrastructure is not just about using graphical user interfaces. It also acts as an abstraction layer that hides the complexities and variability between vendors and data-centers. This allows developers to incorporate infrastructure configuration into the build and release process of the application (called “infrastructure-as-code”).

This reduces the time-to-market for new features, reduces release risk and improves SLAs.

For example, by leveraging a CI/CD-based release process and infrastructure abstraction⁸, Target® was able to grow by more than \$15 billion in 2020, which was greater than the company’s total sales growth over the prior 11 years⁷.

Software defined bare metal vs Converged vs Hyper-converged?

Converged infrastructure typically involves a pre-configured package of software and hardware in a single unit. This enables simplified procurement. A converged infrastructure is typically designed for specific application or workload and while the compute, storage and networking components are physically integrated,

the management of those resources remains the same as siloed resources. They are not “software defined” or only partially.

Hyper-converged infrastructures add another layer of abstraction through the use of virtualization. In a hyper-converged infrastructure the compute, storage and network components are implemented through hypervisors. The infrastructure is software defined but virtual: the network fabric is implemented with virtual switches and the compute units are VMs.

Similar to hyper-converged infrastructure, software defined bare-metal infrastructure or composable infrastructure allows users to easily shift resources around workloads but with physical hardware instead of virtualization.

Unlike a hyper-converged solution that typically doesn't scale beyond 20 or 30 nodes a software defined bare metal infrastructure scales to many data-centers.

Conclusion

Software defined bare metal infrastructure offers strong business benefits.

Software defined infrastructure enables service providers to compete in a ultra competitive, post-cloud business environment where service delivery speed, cost effectiveness and developer-centric features are the norm.

New architectures based on containers will increase the need for bare metal while - at the same time - companies will use a mix of cloud and on-premises, bare metal and virtualized infrastructures making bare metal automation increasingly important.

Why choose MetalSoft

Perhaps the most important factor to consider in your bare metal automation purchase decision is the vendor as well as the overall technical solution's architecture.

The team that built MetalSoft has pioneered Bare Metal-as-a-Service at Bigstep, an UK based bare metal cloud provider focused on Big Data and HPC since 2013.

MetalSoft is:

- 1. Non-proprietary.** Unlike other Software defined bare metal infrastructure solutions on the market, MetalSoft is a purely software, vendor agnostic solution. It leverages industry standards such as eVPN, RedFish, Restconf, Netconf, iSCSI instead of proprietary protocols and components.
- 2. Flexible.** MetalSoft offers more choices in how components are used by having very good network provisioning and local storage provisioning support, specifically designed for high-speed databases, large scale ML/AI that can be used alongside our disaggregated storage support with containers.
- 3. Highly Scalable.** MetalSoft scales to thousands of data-centers with its agent-based architecture, supporting both large scale service provider operations as well as edge scenarios.
- 4. Cost-effective.** MetalSoft offers the benefits of software defined infrastructure but without the cost and vendor lock-in of other solutions. It allows the client to use off-the-shelf, general purpose hardware and turn it into a component of an integrated software defined

infrastructure.

Sources

- ¹ MetalSoft Whitepaper - *ROI of bare metal automation*
- ² Flexera - *State of the cloud 2021 report*
- ³ Research-and Markets - *Bare Metal Cloud - Global Market Trajectory & Analytics*
- ⁴ Research-and Markets - *Bare Metal Cloud - Global Market Trajectory & Analytics*
- ⁵ MetalSoft Whitepaper - *How much time it takes to deploy a server*
- ⁶ MetalSoft Whitepaper - *Performance evaluation of container network latency*
- ⁷ Target Corporation - *Reports Fourth Quarter and Full-Year 2020 Earnings*
- ⁸ Protocol - *Target CIO Mike McNamara makes a cloud declaration of independence*
- ⁹ MetalSoft Whitepaper - *Poll on OpenStack use*
- ¹⁰ Uptime institute - *A Tale of Two Data Centers: Contrast in efficiency*
- ¹¹ Nextplatform - *Relentless competition drives down ethernet switching costs*

Evaluate or Purchase MetalSoft

To test drive MetalSoft today or to find out more contact us at: inquires@metalsoft.io.

Visit <https://www.metalsoft.io/how-it-works> for in-depth technical documentation and additional resources.

