



SMART GUIDE

Essential considerations for optimizing cloud computing costs

Cutting through cloud complexity.



Introduction

It all looked so clear when companies first made the move to the cloud. Superficially, it all made sense. Organizations no longer had to pay for compute power that they didn't use, so it seemed a no-brainer that there would be money to be saved.

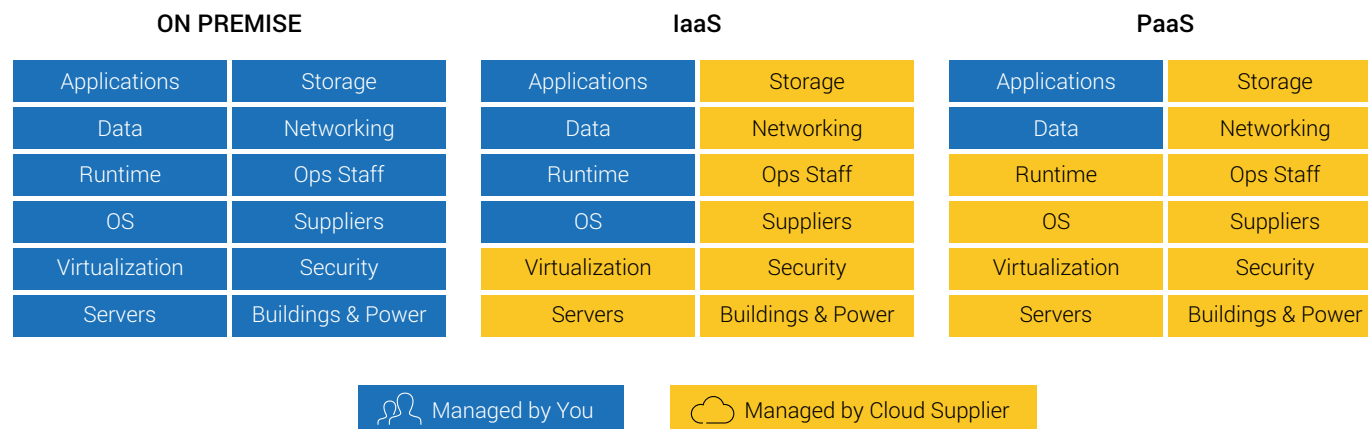
But it didn't work out like that; all too frequently organizations found that the promised savings didn't materialize and that, in many cases, their cloud deployments were actually costing them more than their previous on-premise implementations.

Add to this the additional complexity of how to deploy in the cloud – currently popular models are Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) – both have pro's and con's and understanding these when design decisions are made can greatly affect downstream costs and the effort required to control them. An overview of the typical separation of key responsibilities (and potential sources of cost) is shown below.

Companies are certainly worried that they're not using cloud effectively. According to a survey from 451 Research, more than half of large enterprises, 57 percent, worried about cloud costs on a daily basis and a whopping 80 percent thought that poor management of their cloud deployments was proving to be a financial drain on their companies.

That's a surprisingly large number of people concerned about the management of cloud costs, particularly when there are measures available to keep costs low. Industry research reveals that, perhaps most surprisingly, they're reluctant to take measures such as shutting down unused – or little used – workloads, rightsizing instances or moving from PaaS to IaaS if conditions demand. There could be many reasons for this: ignorance of how much cloud capacity is needed; a lack of knowledge of what's actually out there (something that could be exacerbated by a shortage of skilled cloud practitioners) or, perhaps, badly negotiated contracts.

This Guide explores many of the reasons as to why organizations have had difficulty achieving the anticipated cost savings and discusses opportunities to adopt best practices that help rationalize and optimize cloud deployments to ensure organizations maximize their ROI through Cloud efficiency programs.



Exploring cloud migration cost considerations

In other Scalable guides and blogs we have discussed the pros and cons of the cloud, cloud migration, and explored the importance of establishing a cost baseline as the first step of any cloud migration project.

The basic financial drivers for a cloud migration of an existing business system are usually cost, clarity and predictability and/or cost reduction. That's not to say these are the only drivers; agility is oft quoted, but if cost is your focus then it's worth exploring some aspects to consider.

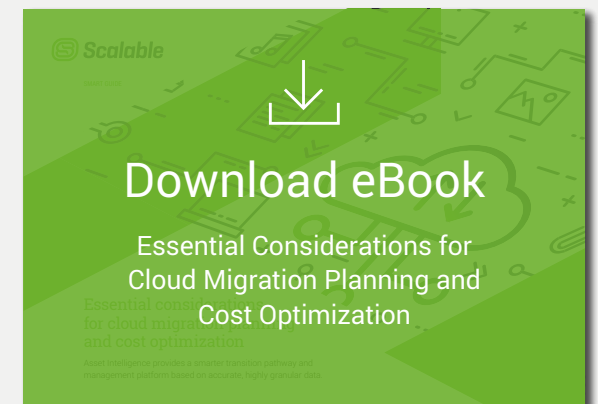
Cloud costs tend to be clear cut and relatively easy to understand, at least statically, before you start growing, but that's for another article. On-premise costs are often shrouded in equipment depreciation, warranty and licensing renewals, maintenance costs, shared staffing and facility costs, management charges, time to remove exiting on-premise costs, etc., which makes them hard to clearly articulate and accurately calculate.

Just gaining an understanding of the total costs of owning and running an existing business service are sometimes enough to justify a move to the cloud.

However, this is a move that requires careful preparation, especially in terms of deep-dive analysis of your existing IT estate and the way you use the assets. If you don't invest enough time in preparation and planning, you will struggle to realize the full benefits.

For more in-depth information download the eBook, "**Essential Considerations for Cloud Migration Planning and Cost Optimization**," to discover best practices for cloud migration and the 6 questions that need to be addressed to ensure success during and after your cloud migration.

- ① What have I got?
- ② What should I move?
- ③ Who can give me what I need?
- ④ How do I move?
- ⑤ How can I keep delivering value?
- ⑥ How can I avoid cloud sprawl and cloud shock?



Optimizing cloud costs in production: minimizing risks to maximize rewards

Trouble in paradise: Cloud Costs – minimize risks & maximize rewards

As enterprise cloud services have become more prevalent, more embedded in core operations, and more diverse, the key challenges have shifted from the hygiene factors such as security, privacy, and sovereignty, that have largely been addressed, and instead they have been replaced with new challenges related to business value realization beginning to emerge as enterprises start measuring the business impact of cloud services.

All that's great...Can amplify all that's bad...

The benefits of cloud, when properly implemented, are well known and attractive: real-time & dynamic on-demand purchase that enables scalability and agility. No need to project future demand and risk overprovisioning to enable economies when purchasing to meet future demand requirements or cope with erratic peak periods of load, and an elimination of redundancy and overprovisioning.

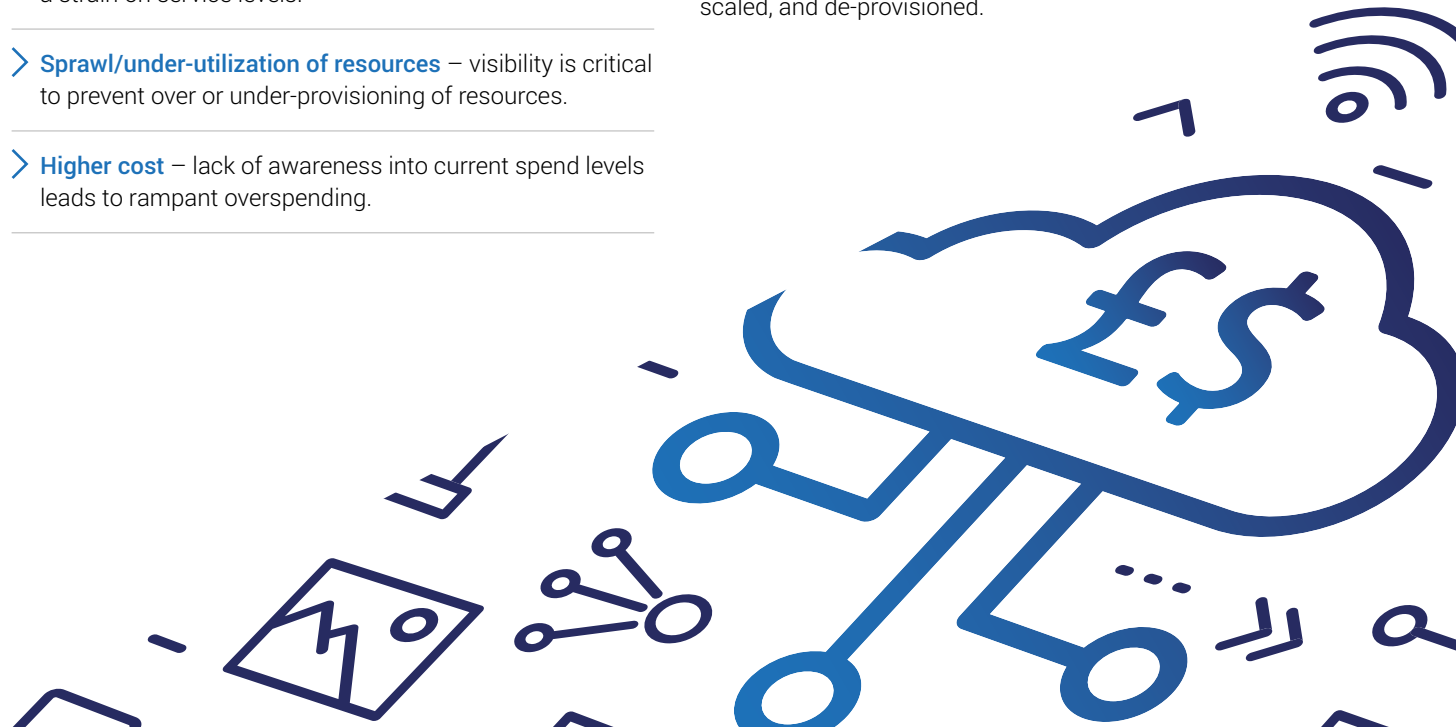
However, the on-demand nature of cloud use can often accelerate and amplify the escalation of cloud costs. Many organizations are exposing themselves to significant risks as a result of a lack of action in adopting processes and technology to monitor and optimize their cloud environments leaving themselves open to financial and operational risks that impact business performance.

Poor Cloud Financial Management can have a serious Negative Impact on the Business

- > **Slower cloud adoption** – without a process for managing how cloud resources are consumed, controlled and budgeted, enterprise adoption slows or comes to a halt.
- > **Negative impact on innovation** – the promise of cloud and freedom to innovate is at risk for organizations who are unable to properly manage decision making at the design stage and escalating cloud spend.
- > **Lower quality of service** – poor utilization of cloud puts a strain on service levels.
- > **Sprawl/under-utilization of resources** – visibility is critical to prevent over or under-provisioning of resources.
- > **Higher cost** – lack of awareness into current spend levels leads to rampant overspending.

Often organizations either ignore the problem, hiding behind the misperception that the amount of overspend is too small and not worth resolving, or there is a reluctance to admit to wasted spend factors, sometimes into the millions, in lost savings and deepening levels of inefficiency. As the growth has occurred organically in many areas there is no central control for spend.

Achieving the goal of efficient cloud use must be more than a one-time event or once-a-quarter focus and organizations need build best practices for continuous optimisation. IT teams and finance staff must become proficient in continuous processes to monitor and optimize cloud spend and use intelligent automation to deliver the analysis and insight needed to efficiently govern costs when new resources are being continually provisioned, scaled, and de-provisioned.




Multi-cloud computing introduces some significant cost management challenges

One of the bits of wisdom that became widely accepted early in the cloud computing era was that an organization that uses multiple providers should be able to save money – in effect, pick the most cost-effective way to deploy its workloads. In reality, the opposite is often true.

A variety of factors makes the reality fall far short of the hope. Even so, many businesses have embraced multi-cloud computing, and, experts say, the right mix of techniques and tools can help rein in costs to some degree and make multi-cloud more workable.

In research from Enterprise Management Associates, 42% of survey respondents listed the desire to control cloud costs as their top IT operations priority. Further, research from analyst firm Enterprise Strategy Group showed that the biggest challenge associated with multi-cloud computing was simply trying to understand what they had running where and how much it cost them.



42% of survey respondents listed the desire to control cloud costs as their top IT operations priority.



Why cloud computing suddenly seems so hard and expensive

Complexity of cloud pricing options – it's not easy to compare

AWS, Microsoft and Google all have different price models and, what's more, they frequently change as the providers try to retain commercial advantages. While cloud pricing can seem simple on the surface – a by minute or hourly cost for a cloud instance or a cost per GB-month for storage – the reality is that there is a dizzying array of options from which to choose. Add to this the challenges of choosing IaaS or PaaS, and there are tens of thousands of pricing options across the three leading cloud providers. Instances can also have significant price differences based on the region in which they run. With IaaS, older versions of instance families can be more expensive than their replacements. Even instances with similar amounts of CPU power and memory can vary markedly based on other add-on characteristics.

Then there's the workload issue to consider – PaaS can be cheaper until workloads become more and more persistent to a point where IaaS becomes a better option.

Organizations who are using cloud will have to stay on their toes to make sure that they're conforming to the current best deals – not easy to achieve and impossible with some tools that only show billing for one provider.

Difficulty selecting the appropriate instance sizes

If going down the IaaS route, as engineers and IT staff build and deploy applications, they need to decide which instances and sizes to provision. In many cases they may be unfamiliar with the performance characteristics of the cloud instances or of the applications they're deploying. When migrating instances from on-premises infrastructure, they may not know what the equivalent instance sizes would be. They will often take a "better safe than sorry" approach and select a larger size. Once the infrastructure is overprovisioned, it rarely gets downsized. Resource owners don't have full visibility into cost implications, and a lack of automation to optimize workload contributes to the problem.

Specific price points are also a concern for moving to the cloud, i.e. are you using the right type of instance – reserved or spot instance. The only way that you can know this is to understand what the performance and availability requirements are. For example if a server is constantly used for over 3 months then it could be a candidate for a reserved instance. An irregular workload which isn't time bound could be a good candidate for a spot instance.

One of the biggest savings that companies make when reviewing the performance data from their cloud instances is when it is being used. Systems such as development or test can be spun down and up when needed, making for large savings.


Again, without consistent monitoring you won't have the information to make these decisions.

A lack of insight from vendor billings

There is no useful middle ground in vendor billings with most providing a high level summary to deep dive granular detail and no actionable insights in between.

Cloud providers' billing tools provide information about the assets rented from them. What they won't tell you is exactly how they're being used and won't provide any insight, for example, on how a company could balance workloads.

A good comparison is with mobile phone bills. Billing data is much like a cell phone bill – either a very high-level summary or the deep dive minute by minute breakdown of everything that's been used and hence paid for. The detail is all there, but you need to analyze it to get what you need from it. It's difficult, time consuming, and often impossible to get insight on how a company is actually using the tools that run on cloud with pinpoint data from an application, including when and how intensively it is used. In addition, vendor bills also don't have any split between cost centers or other company cost charging, which may make cross-charging extremely complex.



Why cloud computing suddenly seems so hard and expensive

Resource owners don't have full visibility into cost implications

At the time of provisioning, resource owners often have little to no visibility into what their applications will cost in the cloud. The hourly cost of cloud instances can seem low, so they may not understand the full impact when they run instances for weeks, months, or years. This limited visibility into costs can also be exacerbated during agile development processes, when teams are automatically provisioning and tearing down deployments for development and QA. If templates or automated scripts are used, instances can be repeatedly overprovisioned, resulting in a continual recurrence of wasted costs. And once applications are running, resource owners may not receive reporting that enables them to see the cost implications of this overprovisioning.

Resource owners and IT don't have visibility of all billable resources and their utilization

Flexible and agile scaling brings a need to continually rightsize: a need to dynamically monitor and respond to waste.

Cloud is a new way of doing IT, and just using it as a replacement for an internal data center could be the wrong decision. The cost shock that companies can receive is often a useful driver for cost base analysis against re-architecting their current services to be designed with cloud in mind.

Without resource owners being able to have easy to access and understand billing they can't understand their general cost benchmark and understand the drivers for an increase or decrease in cost. Resource owners normally have a tolerance of a certain percent of cost increase/decrease but outside of that need to be alerted straight away to review and take the best course of action.

Lack of insight to optimize workloads

Optimization is an ongoing challenge in the enterprise. Even after waste is identified and resolved, the dynamic nature of cloud use means that waste reoccurs. Automation is critical to dynamically monitor and respond to waste. For cloud governance teams to ensure cost-efficient cloud use, they need automated tools that work across all of their cloud resource pools. They need to identify specific areas of waste and collaborate with resource owners to take automated action.



The 5 pillars of cloud cost optimization

Understand the optimal cloud environment

- › Leverage automated data analysis to accurately identify the best-fit cloud configuration and its related cost.
- › Understand the actual usage of the system and what cost benefit are you currently getting

Reduce unnecessary costs

- › Identify idle compute resources and unused storage to avoid cloud sprawl.
- › Shutdown wasteful spending automatically
- › Automate waste reduction by scaling down or halting under-used cloud resources so your teams can focus on innovation.
- › Understand what you are using storage for and are you taking the most cost-efficient approach
- › Ensure you have processes around the archival of data

Plan capacity to support business needs

- › Forget under provisioning or overprovisioning and find the perfect resource for the job.
- › Don't be afraid to refactor the service to reduce costs

Achieve cost clarity and control

- › Understand cloud costs, determine your ideal pricing plan, filter billing data, and build chargeback reports.
- › Ensure understanding of who owns what costs and give them the management tools needed to do this
- › Monitoring of the network is not a one time activity

Utilize data for stronger purchasing decisions

- › Knowledge and insight that delivers a stronger purchasing position
- › Build strategies for reservations and other vendor discounts based off your actual usage so you can have complete confidence that you're paying the lowest possible rate

Cloud cost optimization checklist

Reducing
cloud spend



Instances

- > Remove idle instances
- > Switch off when not in use
- > Review performance metrics



Storage

- > Wrong provision of disk
- > Lack of archive processes



Refactoring

- > Rework application to make best use of cloud technologies
- > Take advantage of lower cost storage



Account Management

- > Assign responsibility for cost management
- > Constant monitoring
- > Granular billing analysis



Conclusion

Cloud computing is the way forward: all companies know this. Every organization will be looking to put some (and in many cases, all) of its infrastructure in the cloud. It's going to improve reliability, make upgrades easier, lead to rapid provisioning of equipment and reduce costs.


At least, that's the theory, but as industry research shows, many managers are fearful that they're paying over the odds for cloud and are not convinced that they're making the most of their opportunities.

It's apparent that a proper audit and thorough assessment of an enterprise estate is the starting point. There's a need to know what's there exactly. Are there instances that are not being used efficiently? Are there long idle periods? Are companies getting the best deal?

Information is essential here. Automated tools can ask the right questions and offer the analysis and insight to help drive efficiencies and optimize costs. The CFO is going to want to know whether a company is getting the best value from its investment in cloud and achieving the optimal ROI that perhaps formed a significant part of the business case.

But it's not just about costs. Yes, it's important to keep control of the bottom line but being agile and responsive to business needs, refocusing resources on innovation, speed of delivery and robustness in the face of adversity are all key objectives but perhaps not as easily measured.

Optimization through consistent improvement processes enabled by technology will provide the visibility, insight and guidance as to whether resources are being deployed in the best way and whether the system is working as efficiently as it should, delivering on the wider goals and benefits provided by Cloud computing.



Asset Vision: Cloud asset intelligence and optimization

Asset Vision enables IT teams to view all cloud usage, including that initiated by business units without IT's involvement and thereby stay in control of overall cloud costs. Accurate granular usage data ensures that smarter IT decisions can be made to optimize cloud investments on an ongoing basis.

Migrating systems to the cloud is one thing – ensuring they are running optimally for the lowest cost is another. The remote, hidden nature of cloud computing makes control of costs something many organizations struggle with. Factors like DevOps teams focussed on deployment targets and timescales, complexity of cloud vendor offerings and pricing, lack of physical visibility of the systems as they grow, lack of accurate analysis tools and a mind shift in how systems are deployed and managed mean many organizations pay way too much. Once an on-premise system has been migrated to a public cloud, it's essential to monitor not just the costs but also the resources being used to support the service. Asset Vision's Cloud Inventory and Cost Analysis capabilities ensure that not only do you have a clear view of all cloud inventory, but that the inventory is being paid for in the most expedient way, and comparisons can even be taken with rival clouds.

Asset Vision supports the three major public clouds – Amazon Web Services, Microsoft Azure and Google Cloud, automatically maintaining a complete view of currently subscribed services. Clear reports ensure you see what you are paying for and who is using it, allowing growth to be controlled and the curtain to be lifted.

Once you know what you have, optimization of cloud spend is a constant process, ensuring cloud assets are being purchased in the most efficient way, and alerting to opportunities to improve. Costs for migrated systems (not just individual parts of migrated systems) can be reported, together with forward planning data for future spend and timing.

The next step on the journey is workload analysis – allowing workload demands and cloud resource demands to be reviewed, and opportunities to downsize and reduce spend identified – all while preserving performance.

Asset Vision provides the raw data, analysis and actionable insights to enable the optimal migration process at the lowest cost – and it does this agentlessly, with nothing to install on critical production business systems.



Asset Vision can support you during cloud migration and beyond

Scalable's Asset Vision platform offers several features that make it ideal for supporting cloud migration and management of the production environment.

Single pane of glass

Many tools only give you the ability to look at those of your assets that reside in the cloud. Scalable's "single pane of glass" concept means that you can see what's going on both in the cloud and in your own on-premise data center, as well as across your workstations and other platforms. This is important when you are planning and managing a cloud migration, and equally so when you've completed your migration and need to make sure it keeps delivering value. You can easily compare different scenarios without needing to worry about which application is running where.

Integrated platform

Many products in this market are modular in a bad sense – they are stitched together from various products (typically one for on-premise and one for cloud), and the components don't always work well together. In addition, each module usually has to be licensed separately. License Scalable's product, on the other hand, and you get a complete, fully integrated solution, designed and built by us from the ground up. You are free to use whichever elements you need from day one.

Normalization

This is the process of translating raw data that the system collects about your IT assets into a meaningful description that helps you with activities such as licensing – and indeed planning what to move to the cloud. Any piece of software or hardware that our software detects gets automatically checked against a central normalization database. This enables us to say that, for example, the task that a particular user is running is part of Microsoft Office 365, and we can then link this fact to the relevant licensing information. Scalable's platform can recognize 80% or more of applications straight out of the box – a strength that greatly reduces the amount of manual effort you will need to put in (though you can of course add your own apps to the catalog – see "Customization" below). This strength is due to our integrated, cloud-based, and crowd-sourced recognition engine, which has more than 1,500,000 entries. In addition to simple recognition and normalization capabilities, we also provide release and end-of-life dates, license metrics, and more – all of which adds up to insights you can act on.

Asset Vision can support you during cloud migration and beyond

Customization and integration

If you need to record asset information unique to your organization, whether manually or automatically, Asset Vision IT asset management products will do the job. Asset Vision provides for custom field creation, either creating entirely new objects or extending existing objects. This feature can be used, for example, to keep track of the locations where an asset is being used, which can help optimize usage. Once defined, customization survives all platform updates and can be incorporated into any of the business intelligence functions.

Highly granular information

Many platforms will simply tell you whether a piece of software is in use at a given location, but our information is much more granular. You get detailed intelligence about how people are interacting with their workstations – for example, you could see that a given user used Outlook for three hours today and that this involved 2,000 keystrokes, and that they later moved on to Excel. The value of this in planning a cloud migration is that you can analyze trends and work out, for example, which users might be satisfied with the browser-based version of Office tools and which need the full-blown application. You could also work out which users might be able to use a cheap Chromebook instead of a high-powered laptop. This helps you go from a high-value capital asset estate to one that uses much cheaper devices or maybe even BYOD.

Full visibility of cloud use.

Asset Vision gives you a complete picture of what cloud apps and services are being used across your organization, including any shadow IT. This makes it possible for management to eliminate cloud sprawl and cloud shock, and to ensure that secure practices are followed. That way, you can keep control without unnecessarily restricting what the business can do.

Intelligent insights, fast.

Asset Vision provides actionable insights after just two weeks of data capture, thanks in part to our inbuilt analysis and reporting engines.



Scalable Software, an innovator in SaaS based Asset Intelligence platforms, provides the critical data and actionable insights that enable organizations to efficiently migrate infrastructure and applications to the cloud, optimize costs and manage compliance.

Worldwide Headquarters

The Future Works, 2 Brunel Way
Slough, SL1 1FQ, United Kingdom

T: +44 203 695 4632
E: sales@scalable.com

US Headquarters

600 Congress Avenue, Suite C100
Austin, TX 78701, United States

T: +1 512-501-2828
E: sales@scalable.com

www.scalable.com

