puppet
Discovery is
the First Step to
Pervasive Automation

Contents

01 Introduction 3

02 Where it all began 5

03 What is discovery? 7

04 Why discovery is more than just monitoring system inventory ..8

05 How Puppet Discovery can help 11

06 Conclusion 14

Introduction

Anyone who's been around the IT and software business for 10 years or more can probably remember a time when counting servers was the same as walking into a data center and counting blinking lights. A "server" was usually a 1U or 2U blade, and each blade was typically dedicated to an application or two running on a flavor of Windows Server or Linux. If you ran IBM hardware, the servers were a little more diversified, but you get the picture.

Of course, virtual machines changed all that and added an order of magnitude to the number of things you had to manage. VMs made it easier than ever to spin up servers, and when public clouds offered similar capability with nearly infinite scale, that added another order of magnitude to the resources you had to keep track of. Containers have made that number climb still higher, into the realm of "I can no longer count these things let alone understand what's on them."

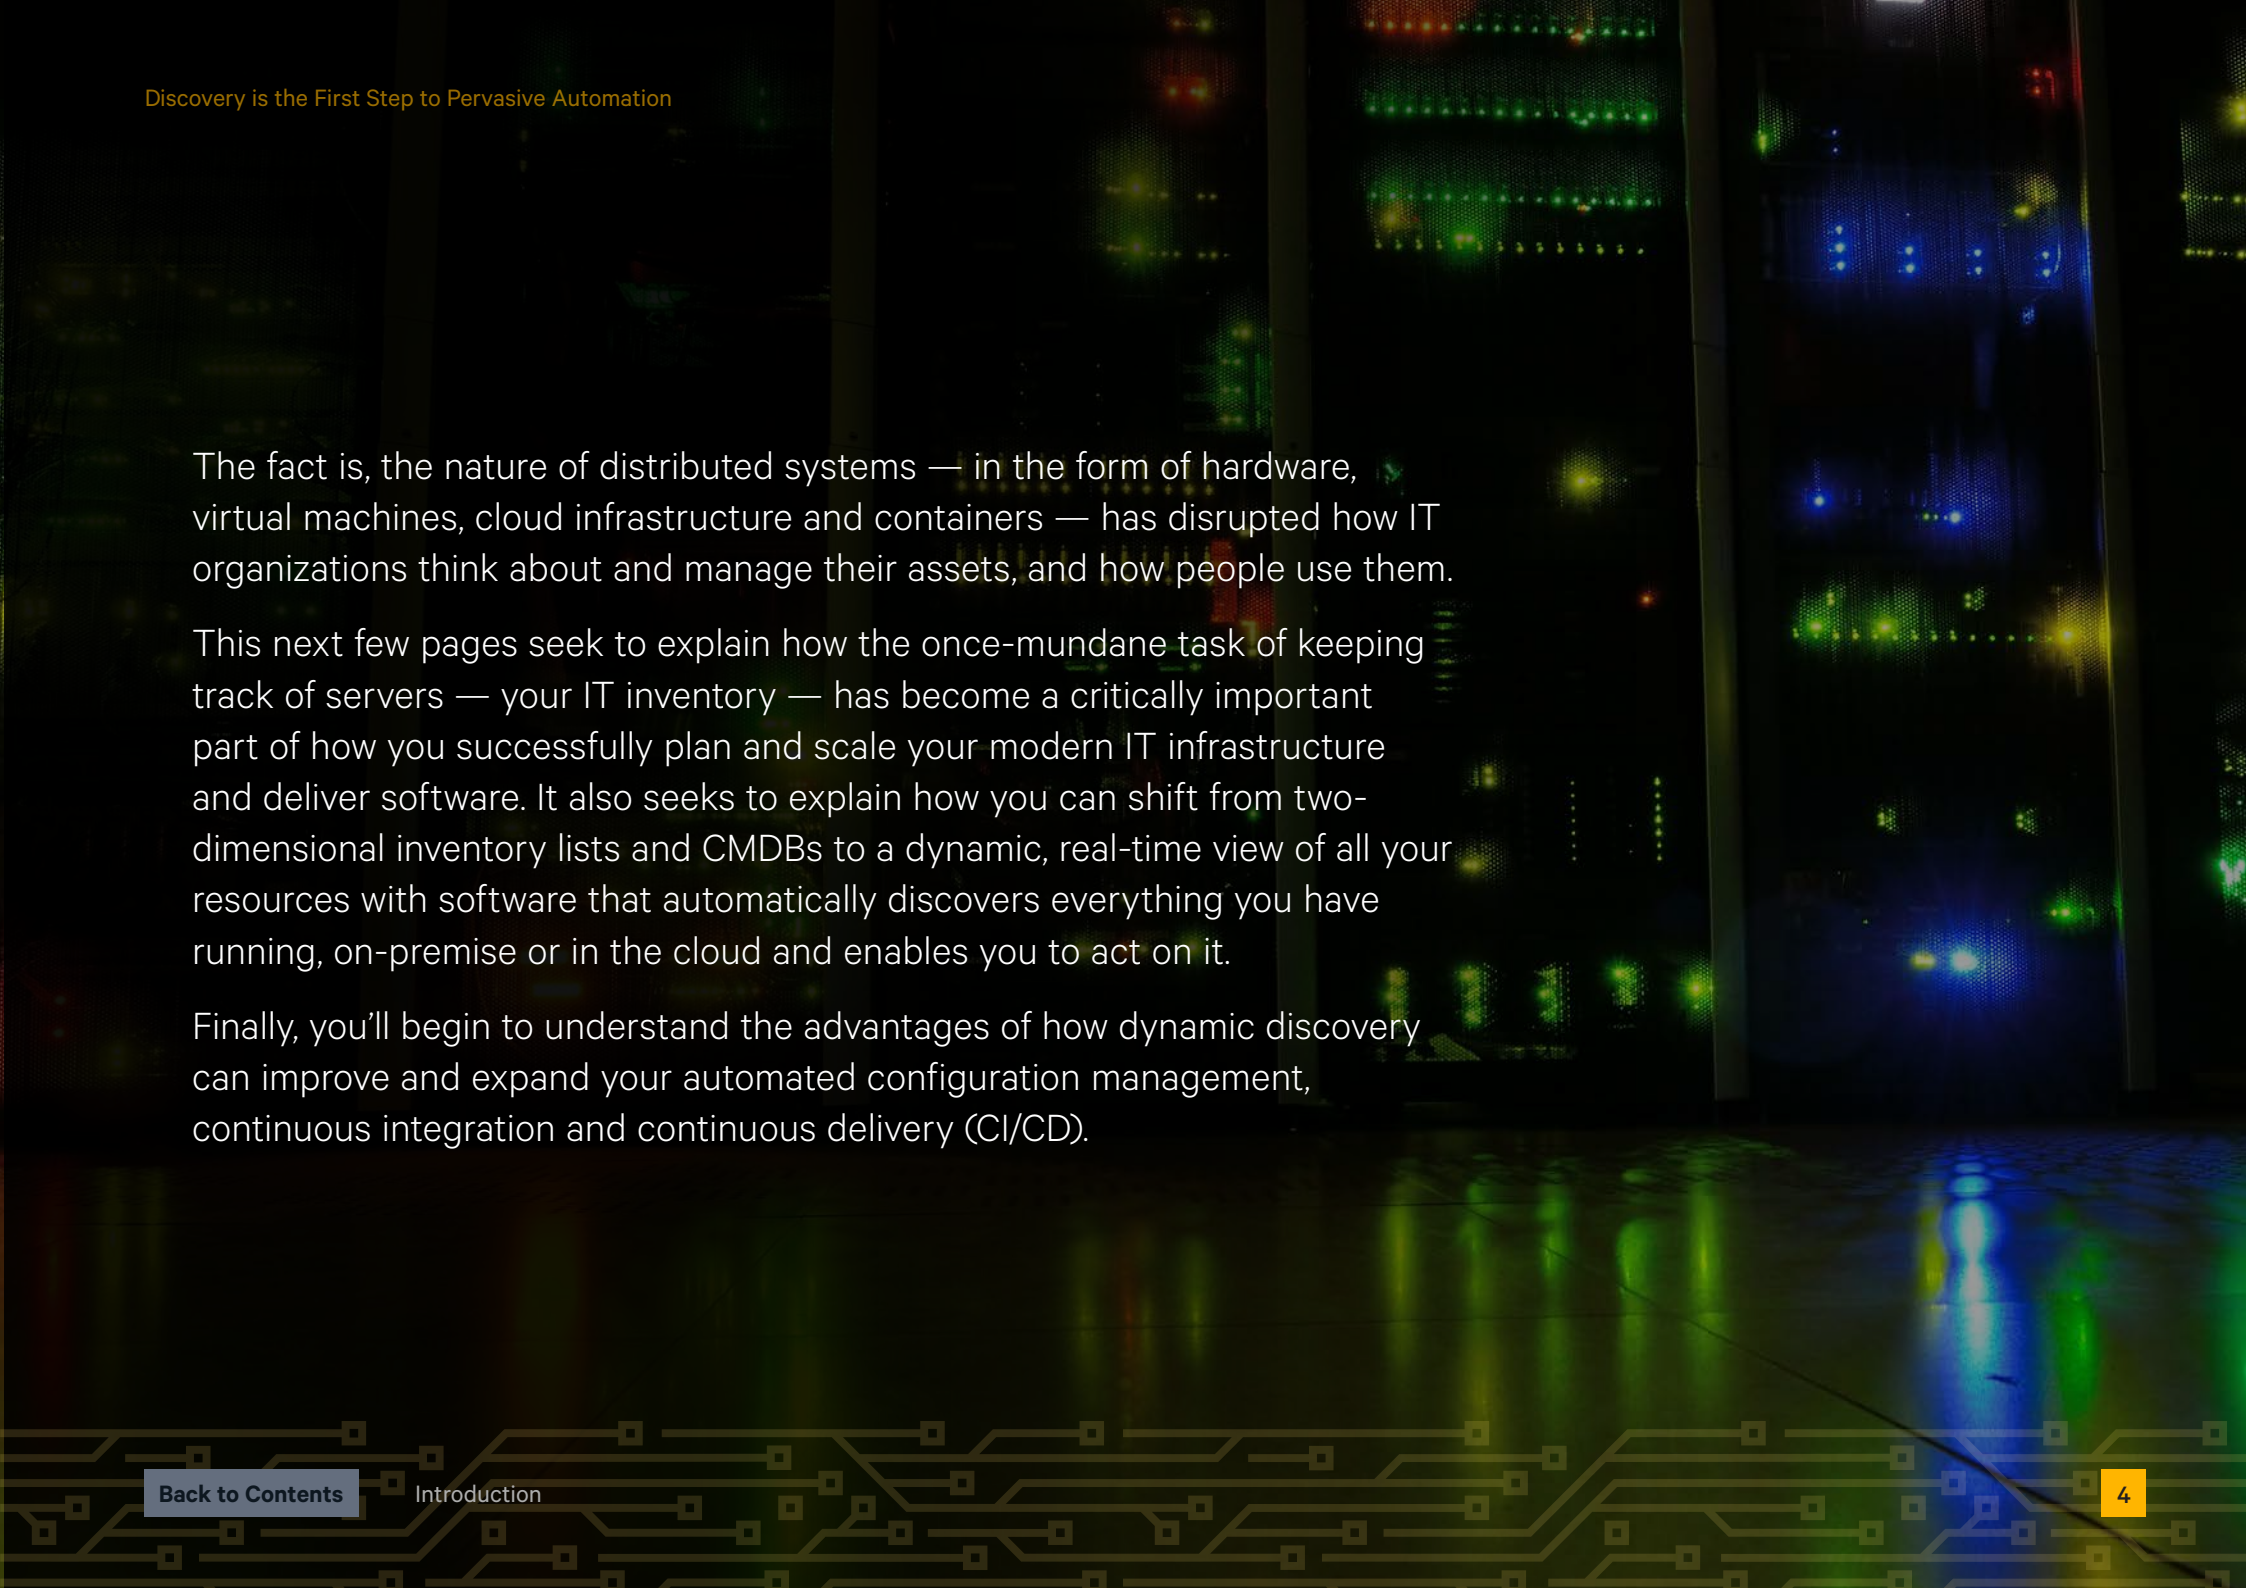

The fact is, the nature of distributed systems — in the form of hardware, virtual machines, cloud infrastructure and containers — has disrupted how IT organizations think about and manage their assets, and how people use them.

This next few pages seek to explain how the once-mundane task of keeping track of servers — your IT inventory — has become a critically important part of how you successfully plan and scale your modern IT infrastructure and deliver software. It also seeks to explain how you can shift from two-dimensional inventory lists and CMDBs to a dynamic, real-time view of all your resources with software that automatically discovers everything you have running, on-premise or in the cloud and enables you to act on it.

Finally, you'll begin to understand the advantages of how dynamic discovery can improve and expand your automated configuration management, continuous integration and continuous delivery (CI/CD).

Where it all began

If you're like most, you and your team probably adapted the company asset-tag system as your first IT inventory system. Perhaps it was a spreadsheet that provided a detailed index of everything from servers and desktops to printers and wireless network devices. The rows of names and types of systems evolved to include IP addresses and, perhaps, OS information, serial numbers and a smattering of other details you all found useful. Maybe you created your own database or portal to do the job, or you advanced to the point where you used a commercial configuration management database (CMDB) that allowed you to better classify your hardware and VMs, and edge away from mere asset inventory to a more meaningful digest.

These solutions definitely have their strengths — simple, easy to share, easy to update — but they start to fall down when you scale, begin moving workloads to the cloud or start adding containers, either in development or production. And truth be told, some CMDBs aren't that easy to set up, and no one ever places "update the inventory spreadsheet" at the top of his priority list.

10.128.1.1	gateway.tucker.com	Cisco	Cisco	Chicago	
DNS Servers					
10.128.1.2	dns01.tucker.com	Ubuntu 16.04 LXC	VMWare	Chicago	1 CPU
10.128.1.3	dns02.tucker.com	Ubuntu 16.04 LXC	VMWare	New York	1 CPU
Physical Machines					
10.128.1.11	esxi01.tucker.com	ESXi 6.5	Del M630	Chicago	16 CPU
10.128.1.12	esxi02.tucker.com	ESXi 6.6	Del M630	Chicago	16 CPU
10.128.1.13	esxi03.tucker.com	ESXi 6.7	Del M630	New York	16 CPU
10.128.1.14	esxi04.tucker.com	ESXi 6.8	Del M630	New York	16 CPU
10.128.1.15	esxi05.tucker.com	ESXi 6.9	Del M630	Seattle	16 CPU
10.128.1.16	esxi06.tucker.com	ESXi 6.10	Del M630	Seattle	16 CPU
10.128.1.17	esxi07.tucker.com	ESXi 6.11	Del M630	London	16 CPU
10.128.1.18	esxi08.tucker.com	ESXi 6.12	Del M630	Tokyo	16 CPU
10.128.1.19	esxi09.tucker.com	ESXi 6.13	Del M630	Sydney	16 CPU
Virtual Machines					
10.128.1.20	mysql.tucker.com	Ubuntu 16.04 LXC	VMWare	Chicago	4 CPU
10.128.1.21	wordpress.tucker.com	Ubuntu 14.04 LXC	LXC	New York	2 CPU
10.128.1.22					
10.128.1.23	lamp.tucker.com	Debian 8.7 (Turnkey)	LXC	New York	2 CPU
10.128.1.24	apache.tucker.com	Ubuntu 16.04 LXC	LXC	London	2 CPU
10.128.1.25	centos.tucker.com	Centos 7 LXC	LXC	London	2 CPU
10.128.1.26	dashboard.tucker.com	Ubuntu 16.04 LXC	LXC	Tokyo	2 CPU
10.128.1.27	windows.tucker.com	Windows 2008 r2	LXC	Seattle	4 CPU
10.128.1.28	windows2012.tucker.com	Windows 2012 r2	VMWare	Seattle	4 CPU
10.128.1.29	debian.tucker.com	Debian 7	VMWare	Chicago	1 CPU
10.128.1.30	puppet.tucker.com	Ubuntu 16.04 LXC	VMWare	Chicago	2 CPU
10.128.1.31	puppet.tucker.com	Ubuntu 16.04 LXC	VMWare	New York	2 CPU
10.128.1.32	elementary.tucker.com	ElementaryOS	VMWare	New York	1 CPU
10.128.1.33	kubernetes.tucker.com	Ubuntu 16.04 LXC	VMWare	New York	4 CPU
10.128.1.34	discovery.tucker.com	Ubuntu 16.04 LXC	VMWare	New York	2 CPU
10.128.1.35	ubuntu.tucker.com	Ubuntu 14.04 LXC	VMWare	Sydney	1 CPU

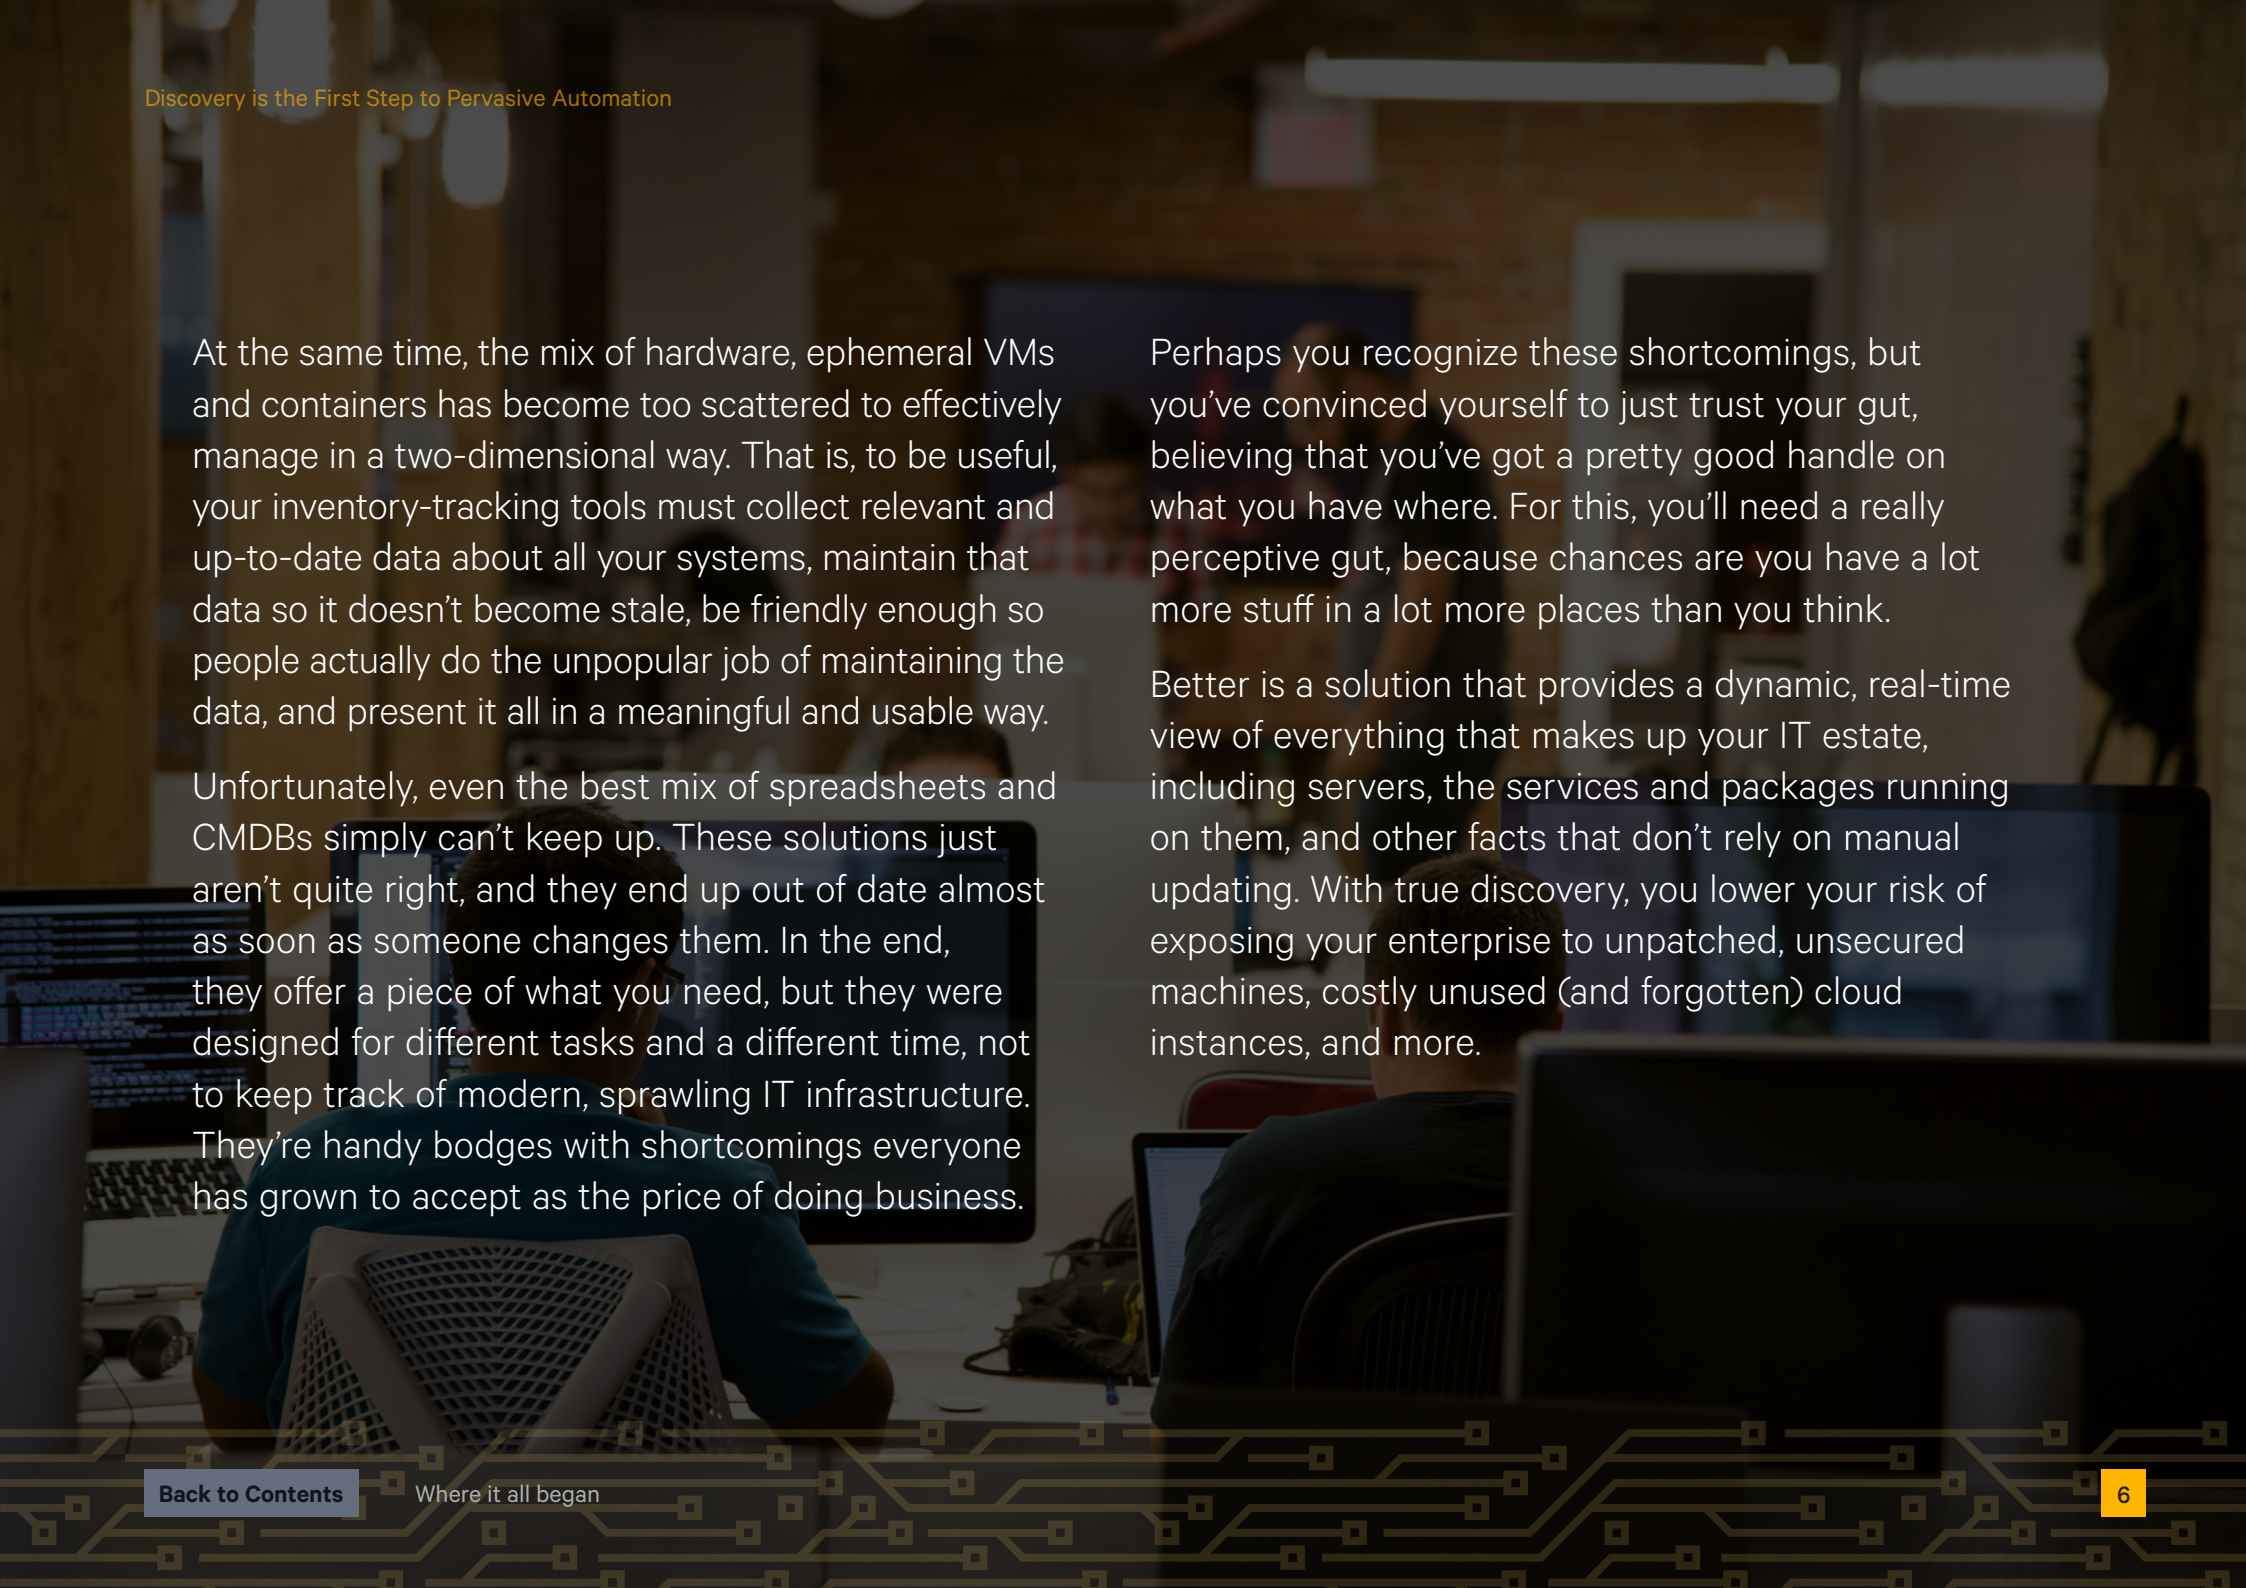

At the same time, the mix of hardware, ephemeral VMs and containers has become too scattered to effectively manage in a two-dimensional way. That is, to be useful, your inventory-tracking tools must collect relevant and up-to-date data about all your systems, maintain that data so it doesn't become stale, be friendly enough so people actually do the unpopular job of maintaining the data, and present it all in a meaningful and usable way.

Unfortunately, even the best mix of spreadsheets and CMDBs simply can't keep up. These solutions just aren't quite right, and they end up out of date almost as soon as someone changes them. In the end, they offer a piece of what you need, but they were designed for different tasks and a different time, not to keep track of modern, sprawling IT infrastructure. They're handy bodges with shortcomings everyone has grown to accept as the price of doing business.

Perhaps you recognize these shortcomings, but you've convinced yourself to just trust your gut, believing that you've got a pretty good handle on what you have where. For this, you'll need a really perceptive gut, because chances are you have a lot more stuff in a lot more places than you think.

Better is a solution that provides a dynamic, real-time view of everything that makes up your IT estate, including servers, the services and packages running on them, and other facts that don't rely on manual updating. With true discovery, you lower your risk of exposing your enterprise to unpatched, unsecured machines, costly unused (and forgotten) cloud instances, and more.

What is discovery?

Discovery, with a lowercase “d,” is all about knowing what you really have, right now, in your on-premise data centers and your public clouds.

Keeping track of your infrastructure in a spreadsheet just doesn’t cut it any more because your organization’s IT infrastructure simply has too many moving parts. Each day brings fresh stuff to lose track of and new opportunities for shadow IT, particularly with the cloud. After all, AWS, Google and all the other cloud vendors make it really easy to plunk down a credit card and spin up an instance.

If you’re committed to the cloud or hybrid environments and multiple cloud providers, you’re even less likely to know what you have at any given moment. That’s why knowing what servers you have, what they’re running, how they’ve changed, why they’ve changed and who changed them when is a critical part of every IT shop’s business and sanity, and the core of automation.

In order to get started though, you need to know what you have before, during and after you deploy it. Discovery helps you get there.

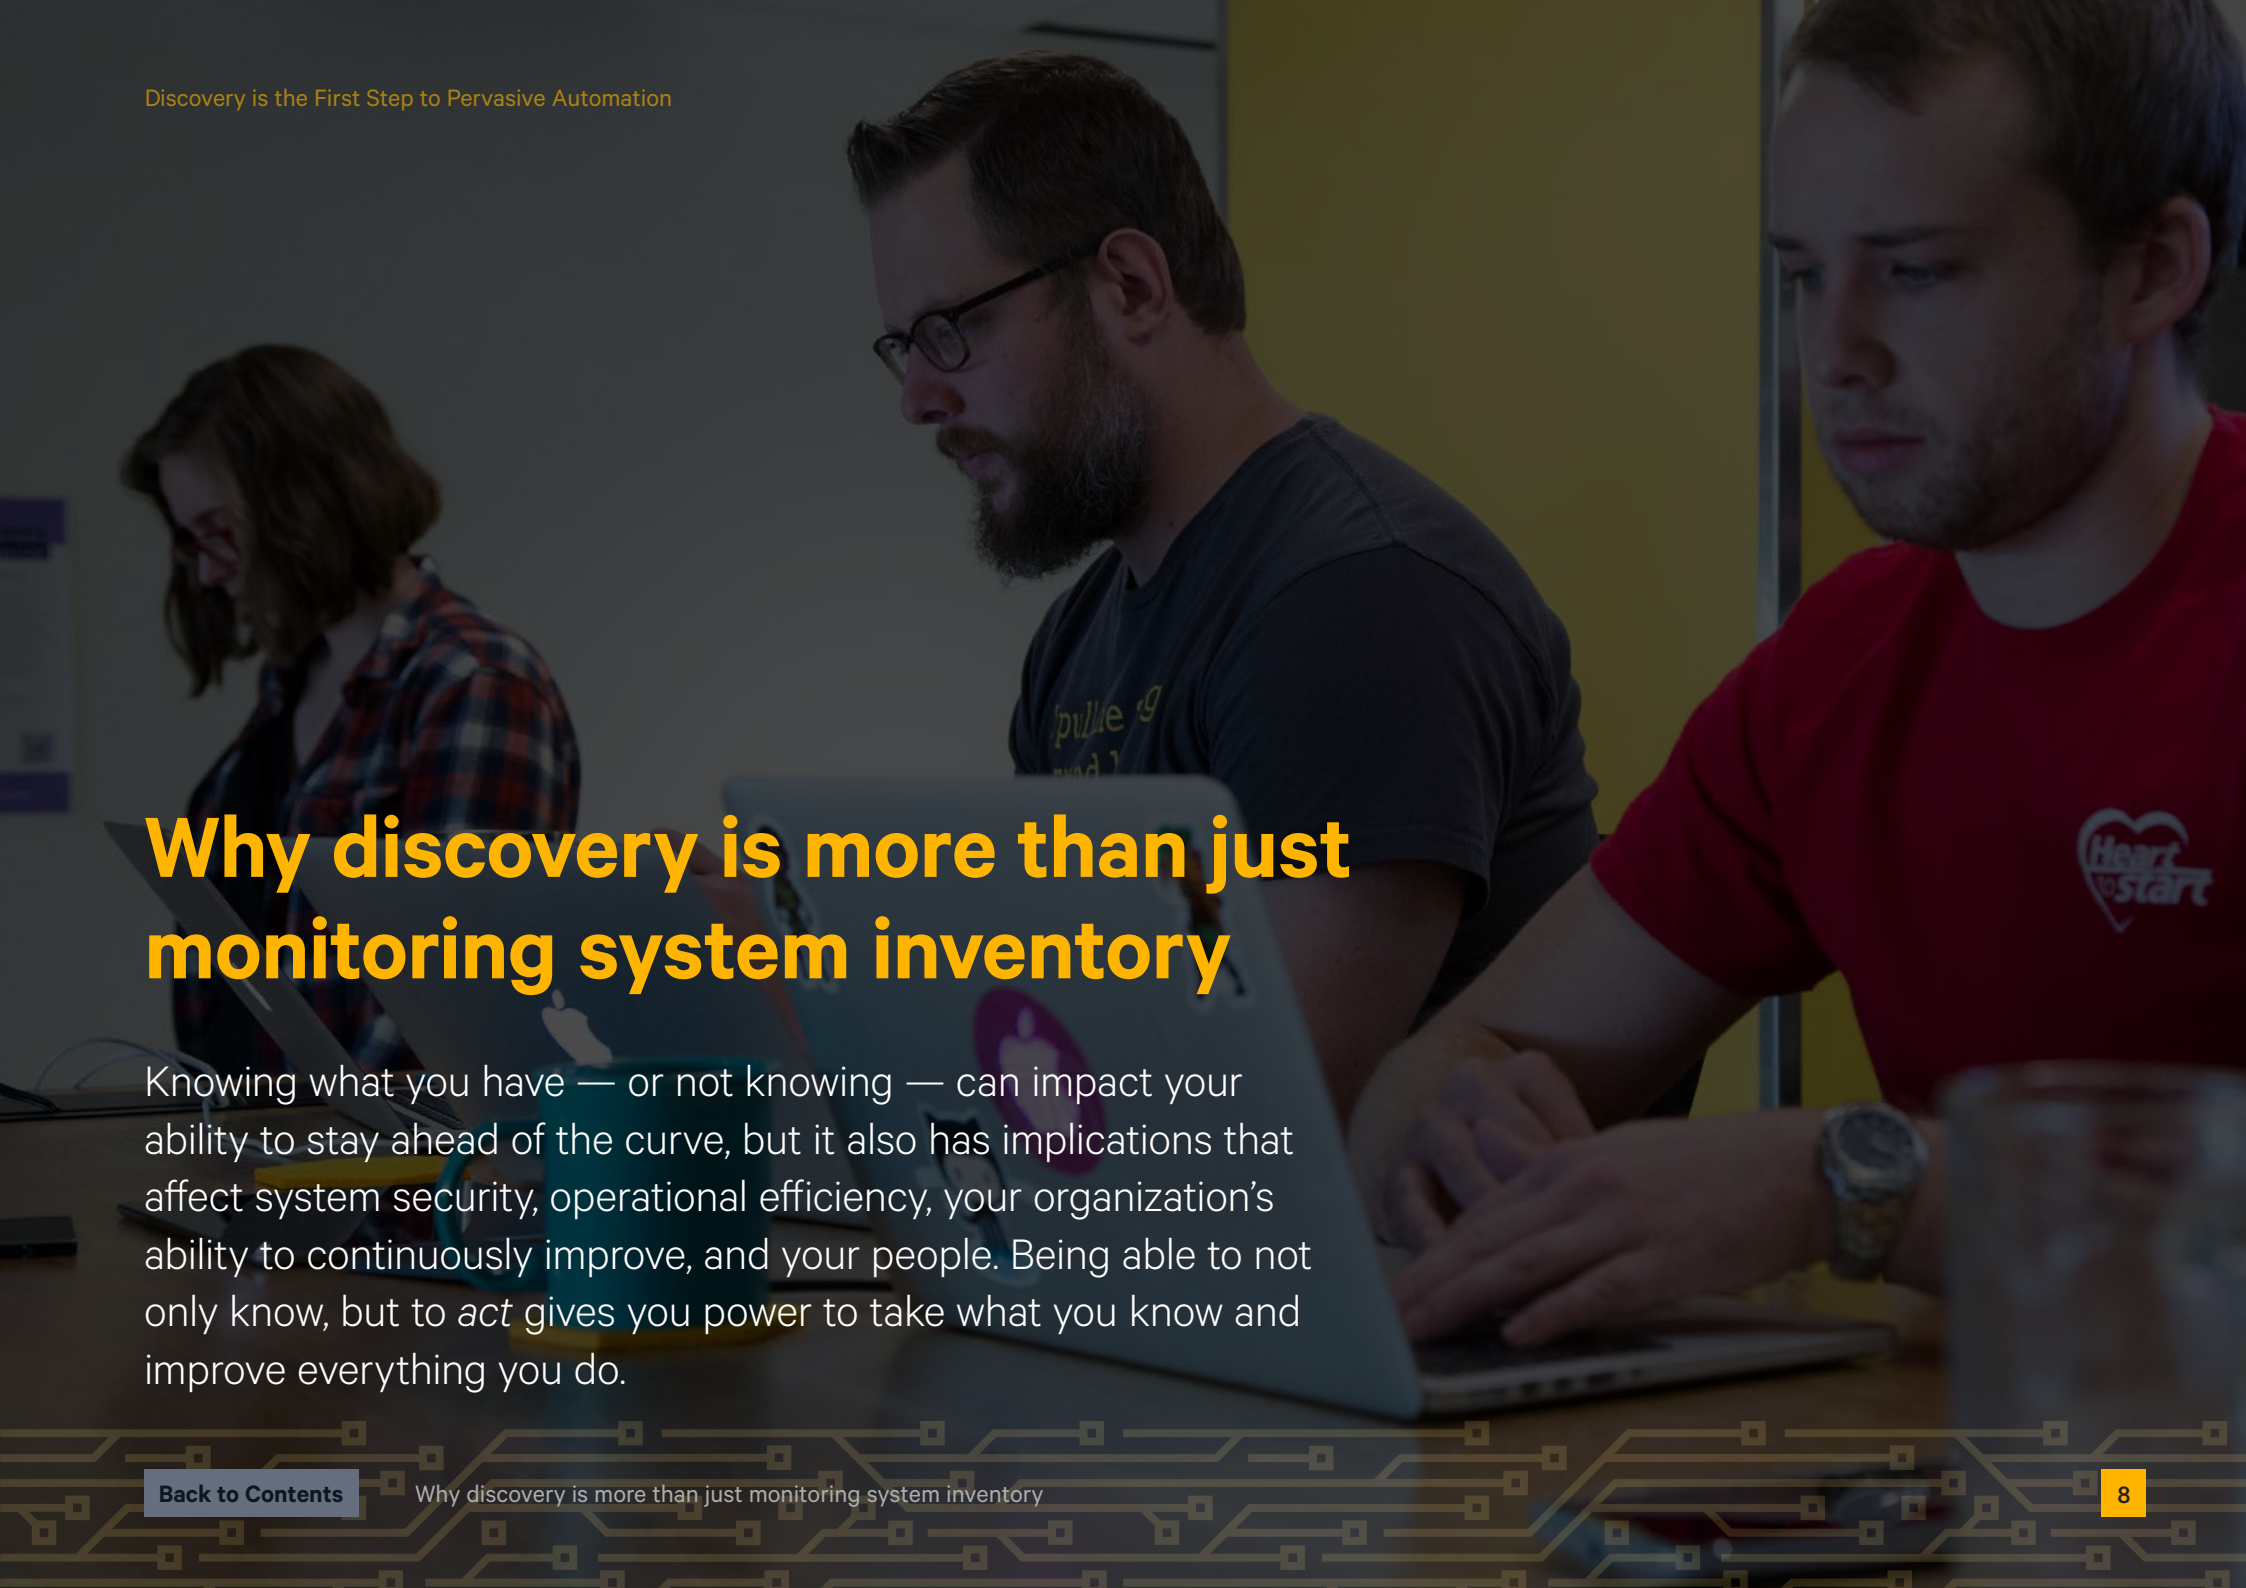

Why discovery is more than just monitoring system inventory

Knowing what you have — or not knowing — can impact your ability to stay ahead of the curve, but it also has implications that affect system security, operational efficiency, your organization's ability to continuously improve, and your people. Being able to not only know, but to *act* gives you power to take what you know and improve everything you do.

Security

This is often the biggest concern for most teams, especially as more data breaches hit the news. You definitely don't want your organization in these headlines. The answer is to think about how you'll enforce your organization's security and compliance policies and procedures across both your on-premise and cloud infrastructure. Your security stance will be assured if you can easily implement in your cloud infrastructure the same tried-and-true policies and procedures that you've already established for your on-premise data center resources.

If you have inconsistent firewall rules, unpatched packages, defunct users with system access, and rogue devices — shadow IT or VMs that were spun up and forgotten about — you have security risks. When it comes to security, knowing what you have isn't just a nice-to-have, it's a must-have.

Operational efficiency

If you're underutilizing your on-premise hardware and VM infrastructure, you're wasting money. Back when a single 1U blade was a single server, the paradigm shifted to take advantage of all that unused compute power and storage. Likewise, if you've virtualized, but you continue to underutilize your VM infrastructure, you're wasting money on everything from hardware to licenses.

At the same time, if you have orphaned nodes in AWS, Google Compute Engine or another cloud, you're needlessly bleeding money. Perhaps a few developers legitimately spun up cloud instances, but no one ever shut them down. That's costing you. Each cloud vendor's dashboard gives you some insight, but if you've spread your infrastructure across several and you're scaling to thousands of nodes, it's easy to lose track — and lose money. If you're trying to do this by hand and juggling different tools to do it, you're bound to miss something and waste money.

Continuous improvement

A big part of running IT or developing software today is improving both the speed and quality of what you do. Competition from born-in-the-cloud companies who do this naturally means you have to improve in order to compete. Your customers expect regular software improvements, faster web experiences, and mobile apps that are smart and intuitive.

For many, automation is the key to improving the things that matter most. But you can't automate what you don't know you have, so instead of having pervasive automation across your organization, you have automation in pockets. If some things are automated and others aren't — maybe all your database and web servers are managed by Puppet Enterprise, but none of your Docker hosts are — you'll spend more time than you need to on mundane tasks and improving the efficiency of what you do. If you don't have a good handle on all your infrastructure, you simply can't enjoy the speed, agility, and stability enabled by pervasive automation.

People

No one likes to fight fires, but teams are often disrupted by unexpected system behavior and unexplained outages. The worst are problems that seemingly "come from nowhere": forgotten or rogue or shadow IT systems with defunct user profiles, non-standard firewall rules, or worse. That constant firefighting makes people miserable, lowers morale, and leads to low performance that keeps your team from innovating, pushing for improvements or looking for ways to make customers happy.

At the same time, no one likes to update a spreadsheet, manage a CMDB or any other cumbersome inventory solution. In the end, people want tools that are easy to use and maintain because they can learn them quickly and start using them right away. The exception are those looking to justify their jobs doing maintenance of older, more cumbersome systems.

Good inventorying and discovery tools help reduce both the number of emergencies and the irksome maintenance tasks, and lower the chance of surprises from the unknown.

Take action

Knowing what you have moves the needle, but knowing alone is only part of the solution to improving what you do. Having the ability to take action — start and stop services, deploy management agents, and know what to automate first or next — turns two-dimensional inventory lists and tools into the jumping-off points for improving and accelerating how you do business.

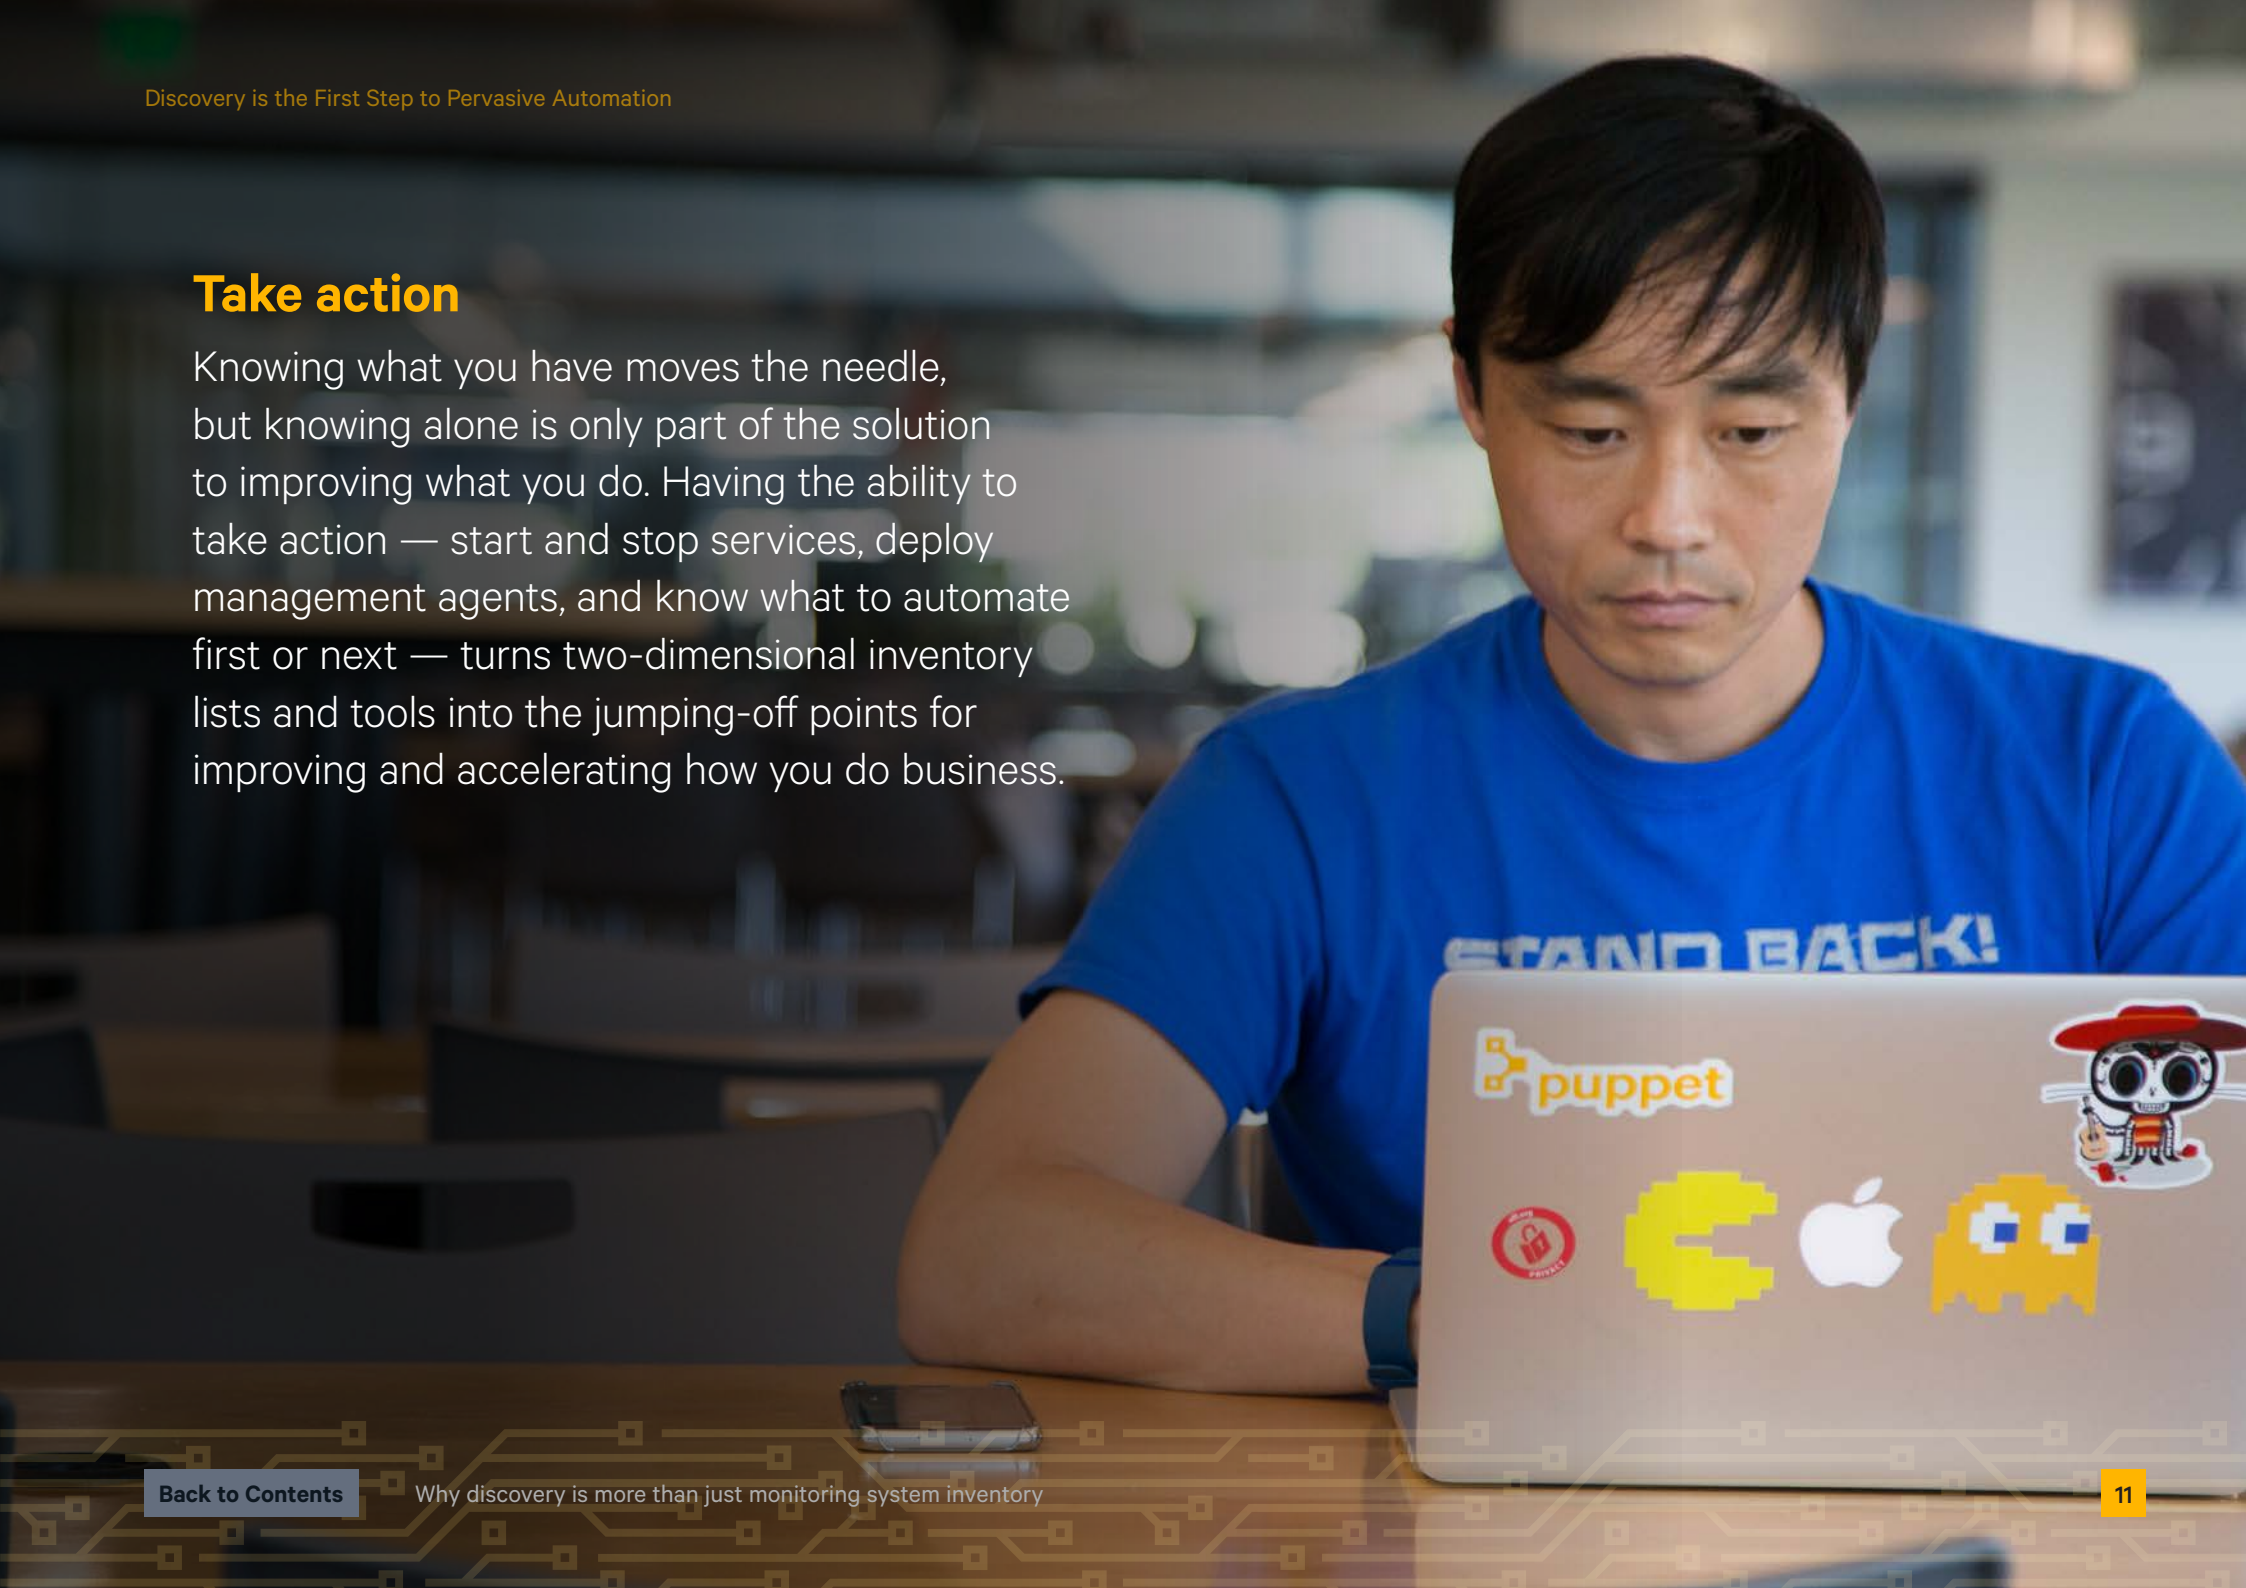

How Puppet Discovery can help

Puppet Discovery gives you insight into your infrastructure, helps you identify the critical elements that should be automated first, and enables you and your team to clearly understand what should be your top priority. It also enables you to act on everything you find.

Puppet Discovery is a stand-alone tool that makes it easy for organizations of any size to dynamically inventory their on-premise, cloud, and container infrastructures, gain insights into discovered resources, and take action to bring them under management. It uses truly agentless, secure connections to discover configuration resources on nodes and present facts about them in a dashboard you can use to deploy Puppet Enterprise agents, stop and restart services, and report on your organization's entire IT estate.

It's the only solution on the market that handles traditional, cloud-native, and containers all as first-class citizens, and enables users to take action on the resources they forage. As a result, Puppet Discovery can help you and your organization manage resource proliferation, reduce vulnerability, and set you on a path toward cloud migration and DevOps.

Unlike a spreadsheet or CMDB, Puppet Discovery is purpose-built for the needs of modern IT and software organizations. Instead of relying on manual updates, it gives you and your team a dynamic, real-time view of your infrastructure.

With Puppet Discovery, you can:

- Quickly discover everything you have, no matter where it lives — on-premise or in the cloud
- Discover Linux, Windows, AWS, Google Cloud and any network-accessible node using just SSH, WinRM or the API credentials you use with your cloud vendors. No agent is required, which means you don't have to find everything before you can inventory and act on it.
- Scan entire subnets or cloud instances so you can quickly see what's living where. You can add endpoints one at a time, too, but if you have multiple subnets and thousands of nodes, you can just point Discovery at them and quickly learn what's out there and what has a pulse.

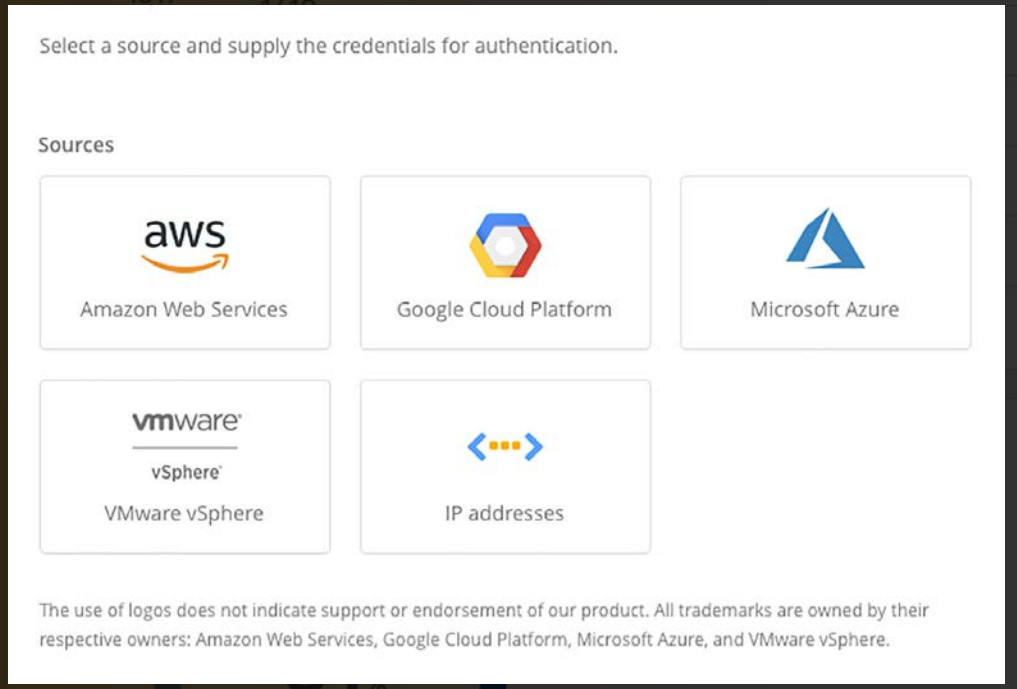

- Take action on discovered nodes and resources. You can stop and restart services, kill defunct cloud instances, deploy Puppet Enterprise agents, deploy third-party agents, and more — all from a single dashboard.
- Quickly get insight into what's running where, including OS types, packages, package versions, containers, and more.
- Quickly bring discovered nodes under management by deploying the Puppet Enterprise agent. This is the key to adopting broader automation and making the advantages of automation pervasive across your organization.
- Set the stage for automation with Puppet Enterprise by discovering what needed to be automated next.

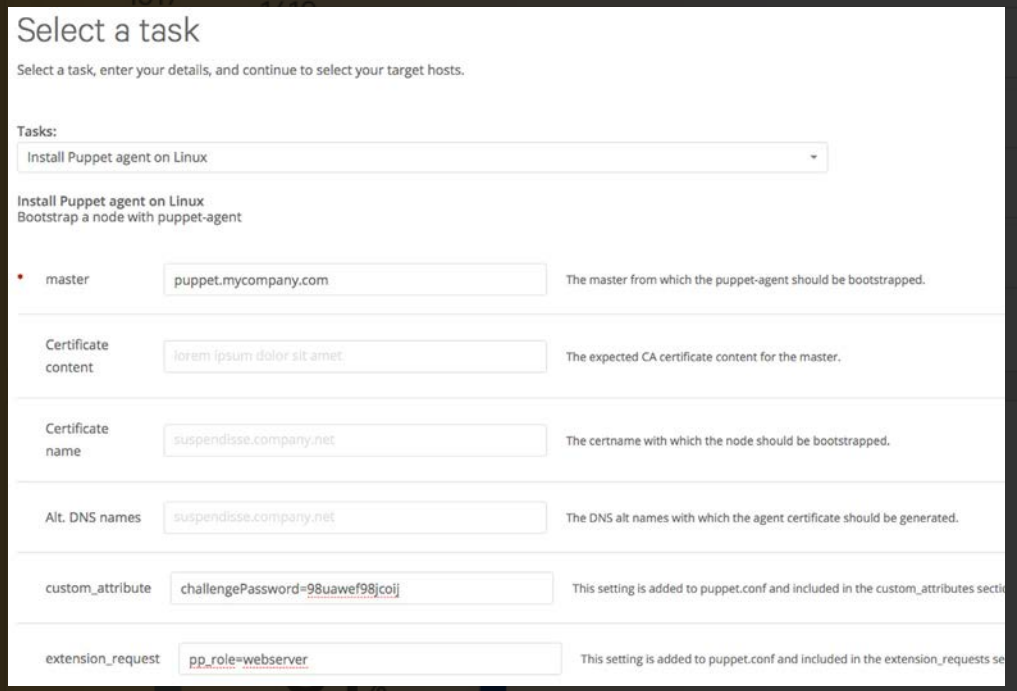

Conclusion

As organizations continue to struggle to manage what they have, it's difficult to understand which parts of your IT infrastructure are vulnerable, draining your budget, eluding your staff, or ripe for automation. Compound that with resource sprawl, particularly in ephemeral environments like cloud and containers, and your work keeps getting harder.

Instead of spending time manually updating spreadsheets, Puppet Discovery lets you see everything you have in real time across your on-premises, cloud and container infrastructure. With full visibility into your infrastructure, plus meaningful insights provided in the dashboard, you can easily get deep insight into your infrastructure and bring it under management.

Unlike other solutions, Puppet Discovery is purpose-built to bridge the gap between initial discovery of your infrastructure resources and automated management, the key to pervasive automation and delivering better software, faster. And it's the only solution that allows you to discover and act on your inventory without switching to another tool.

Download your free trial of Puppet Discovery today and take control of your infrastructure — no matter where it lives.

Resources

Download Puppet Discovery 1.0
puppet.com/discovery

Contact sales
puppet.com/sales

Product Spotlight Webinar: Puppet Discovery
puppet.com/discovery-webinar

Puppet is driving the movement to a world of unconstrained software change. Its revolutionary platform is the industry standard for automating the delivery and operation of the software that powers everything around us. More than 40,000 companies — including more than 75 percent of the Fortune 100 — use Puppet's open source and commercial solutions to adopt DevOps practices, achieve situational awareness and drive software change with confidence. Headquartered in Portland, Oregon, Puppet is a privately held company with more than 500 employees around the world.

Learn more at puppet.com