DCD

> DCD Power Supplement

Powered by

VERTIV™

INSIDE

Why lithium-ion batteries are coming to your facility

> Lead acid has had a long reign, but Li-ion saves space and energy. It could be time to switch

No one dethrones the PDU

> DC power distribution is great in certain circumstances, but the vast majority of data centers still need PDUs to power the racks

Carborundum may grind out eco-mode UPS

> If silicon carbide can reduce the losses in high frequency power silicon, The UPS may change

Powered by

Contents

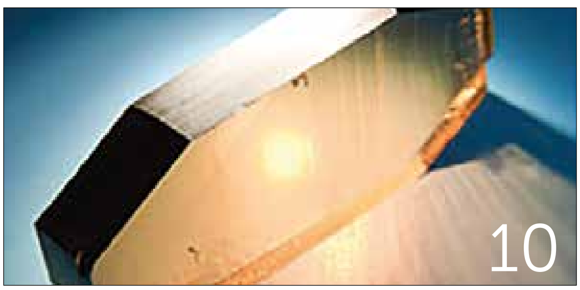

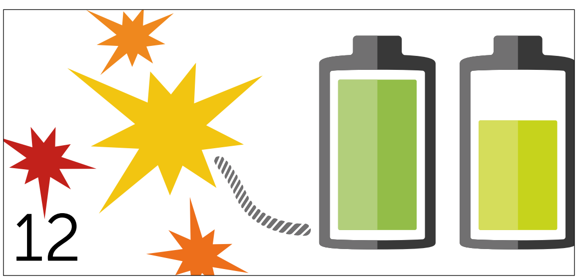

It's all about the chemistry

You may be thinking that the basic science for power distribution is all done and dusted. If you do think that, you would be wrong. It seems like all the leading edge research that will impact on data center power systems is being carried out by people in white coats, using steaming test tubes.

In other words, it's all about the chemistry.

Lithium-ion batteries could revolutionize data center uninterruptible power systems (UPS). They're lighter, longer-lasting and more efficient. Vertiv makes a cogent case for a rapid switch to the new technology (p8). So why isn't it happening immediately?

Part of the reason must be the steady stream of stories about exploding batteries in various consumer devices, like Samsung's Galaxy Note 7.

Of course, the comparison isn't a fair one, it just affects perceptions, and data center people can be very conservative.

In fact, the Lithium-ion industry is expanding rapidly, tracking the growth of the electric car sector. Reliable Li-ion batteries are now available, so this switch is only a matter of time (p12).

Silicon carbide or carborundum is normally used as an abrasive. Professor Ian Bitterlin (p14) believes it could also smooth the transition of UPS systems between the grid and battery power.

The benefit is that silicon carbide (SiC), used instead of plain silicon, can allow currents to be switched quickly with a much lower loss than silicon.

Fast switching is essential in a UPS, and low power loss is an ideal to aim for. So silicon carbide could be an answer to our prayers.

Gallium arsenide (GaN) was once seen as a competitor with silicon carbide in the race to replace silicon in power chips. SiC reached the market first, but the GaN proponents say there's a niche for for their alternative. GaN promises more efficient power conversion, and could allow data center efficiencies to continue improving beyond the end of Moore's Law (p10).

As with many new developments, the field is currently going to early developers. You can't buy this yet, but in not too many years, you may well want to.

PDUs have little to do with chemistry. They distribute power within data center racks, and have done so for some time.

When bodies like the Open Compute Project (OCP) suggest we can shift to DC power distribution (p4), without a PDU, we sit up and take notice. Could data centers all get on board the OCP's bus bar?

We asked the the industry, and it seems that so far the vendors selling PDUs are not exactly quaking in their boots.

The bus bar can really improve efficiency in the right kind of data center - large monolithic web scale facilities. But so far these facilities make up a small proportion of all data centers.

Peter Judge
DCD Global Editor

No one dethrones the PDU

Despite proposals for DC power distribution, the PDU industry sees a healthy future, *Peter Judge* reports

Peter Judge
Global Editor

The Open Compute Project (OCP) was supposed to spell doomsday for power distribution units (PDUs). The Facebook-backed project proposed a DC power distribution system with no PDU in the rack. But the power industry, it seems, is not concerned.

Facebook started the OCP in 2011, planning to spread the use of no-frills customized data center hardware beyond its own webscale facilities to the enterprise.

Large users like Facebook and Google with giant monolithic data centers can afford to have hardware built to their own specifications, dismissing "vanity" features such as badges and cases and using other short cuts.

Through OCP, Facebook shared this approach, and its hardware specifications, as open source, so any organization can use them, including general purpose data centers, run by enterprises and outsourcers.

Ideas within OCP included server designs and switches, but its attack on the fundamentals of data center building created a stir.

Among OCP's flagship ideas is the Open Rack, a wider than

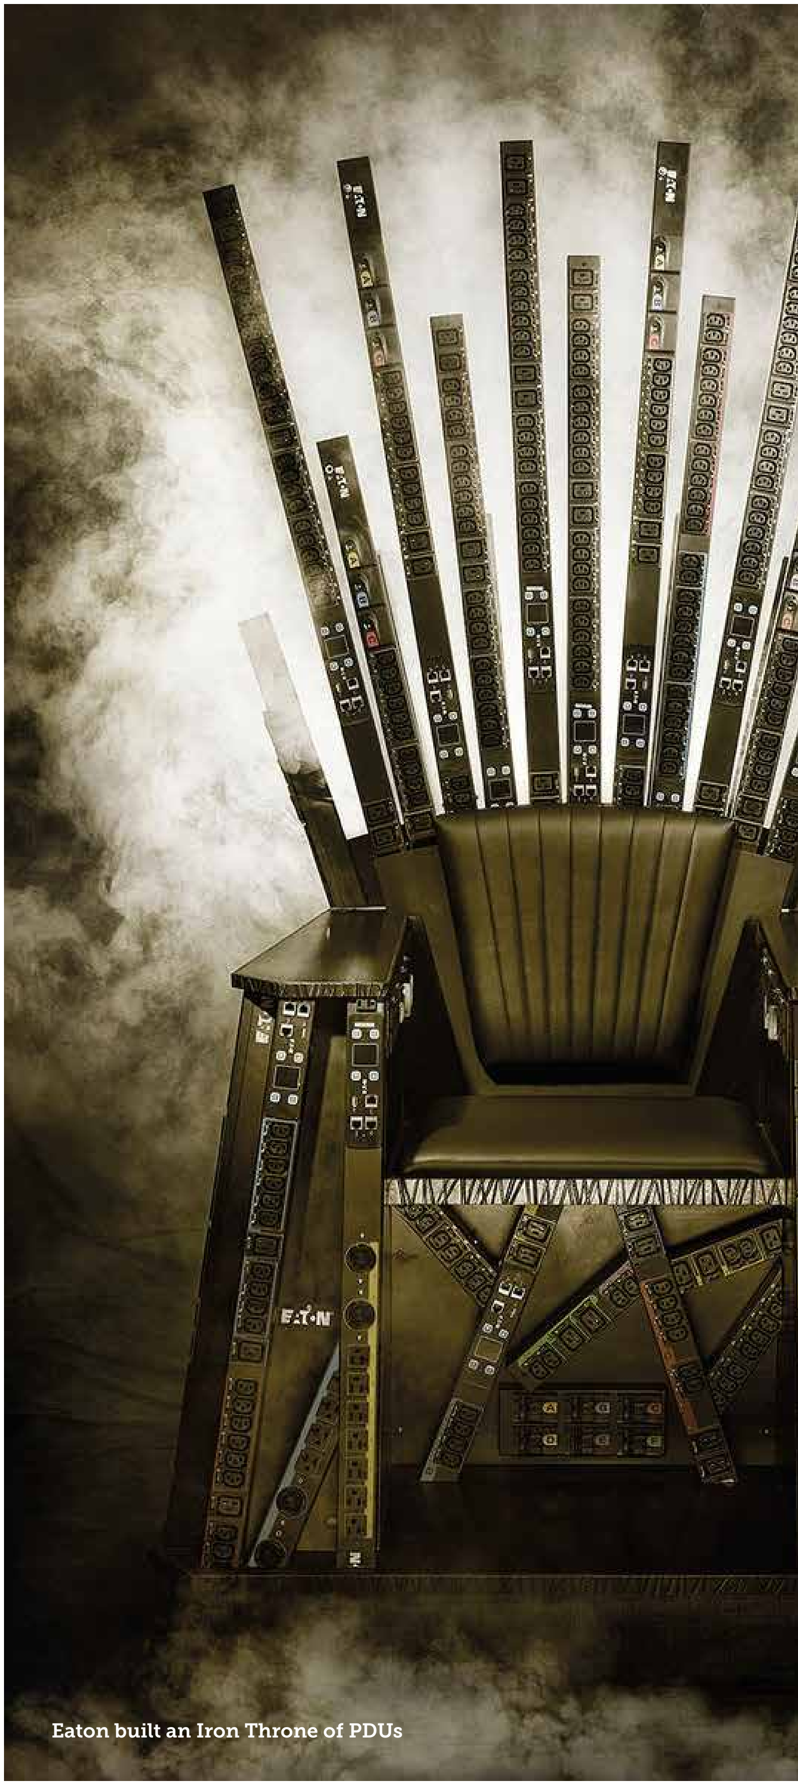

Eaton built an Iron Throne of PDUs

usual 23-inch rack that distributes DC power through a live "bus bar" at the back. This would have no traditional PDU. Other ideas followed, including a normal-sized Open19 19-inch rack from LinkedIn, also powered from a bus bar.

It's immediately obvious that the DC bus bar in the Open Rack and Open19 specifications replaces the power distribution unit or PDU, which delivers AC power through metered outlets in the rack. If OCP took over, said some commentators, it spelled the end of the PDU.

PDU vendors take a more long-term view than that. First of all, the attention given to OCP is out of proportion to its actual market size.

Even if bus bars take off, this is just another step in the evolution of the industry, said Gordon Hutchison, vice president of international operations at Geist: "The OCP approach is relatively new, and may not be embraced by the entire industry. It still requires power distribution, it just changes the nature of it."

In fact, OCP has actually solicited and shared power distribution devices, according to Hutchison's colleague Brad Wilson, president of Geist Global: "These were PDUs in every sense of the word. They have a normal PDU form factor, but just didn't use the connectors we are used to seeing. They had a mix of AC and DC."

Those systems are designed to be easier to customize, and predate the Open Rack with its bus bar, which Wilson says is more of a "radical departure." DC power is distributed directly to servers and other equipment in the rack.

This is more efficient for large homogeneous installations, where every item in the rack is made specially for a customer like Facebook or Google, but very few organizations are built that way. Amazon also works this way, but is bigger still, and able to make its own specifications and keep them hidden.

The people going with Open Rack are the giants who are webscale to an absurd degree - and they still make up a very small proportion of the total racks in the world, vendors told *DCD*. In enterprise data centers, or even cloud providers, the equipment is heterogeneous, and much of it has to be bought off the shelf.

"Not all customers are made alike," said Henry Hsu, vice president of product management at Raritan. "Today, very few of them have the kind of homogeneous data center, made up of standard compute and storage units, that can make use of architectures like Open Rack."

It's tough to estimate, and no vendors could provide a figure, but it seems only a tiny minority, a few percent at most, use bus bars.

<5%
best guess at bus bar market share

Plenty of customers have kicked the tires of the concept, coming along to conferences, or even commissioning bus bar based installations. But in the enterprise these are more or less test rigs.

They have a power shelf to convert between AC and DC, and the customer probably ends up with only one or two cabinets.

In these situations, AC power distribution is cheaper because of the economy of scale. A bus bar in itself is a solid piece of metal, and it is expensive to retool individual equipment to work with it.

Ironically, while the power connectors in conventional power distribution systems are totally standard, the connectors for the bus bar can turn out to be available from a small number of vendors, say PDU players. "As OCP has rolled on, a lot of proprietariness has crept in," one vendor commented.

Even where the connectors are standard, the fact that they are not so widely used as normal power sockets can be a drawback.

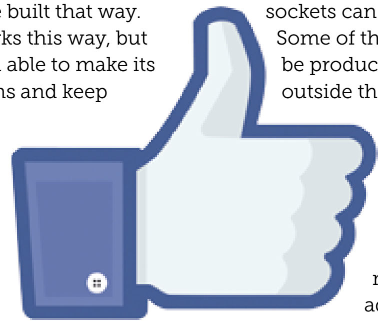

Some of these designs may be produced by people from outside the PDU industry. They can be radically different, but might fall down on issues including manufacturability, adding costs and potential incompatibilities, vendors warned. If they don't work first time, there can be fingerpointing by the different vendors involved.

Conventional PDUs have standards from ANSI, IEC and ISO, which are created by people literally spending 15 years around a table agreeing on them. The downside is that 15 year delay, but the upside is the standards are completely solid.

By contrast, OCP and Open19 are "structured more like Wikipedia," PDU vendors said. They can be modified and evolved, and aren't locked down. This gives flexibility, but can be risky, unless you are a giant player. ▶

What is a bus bar connector?

The OCP (Open Compute Project) specifies 12V bus bars to distribute power up the back of a rack from a power shelf to the so-called Innovation Zone where IT components reside. The specification has gone through two versions, v1 and v2.

The Open Rack v1 design includes connectors from Methode, but the specification says that equipment makers are free to deviate from the exact connectors specified in that design, and use a different clip or method to attach to the bus bar. The responsibility to ensure that the equipment always mates with the bus bar and the rack rests with the designer.

A number of suppliers provide these connectors, including Methode and TE Connectivity. Connectors are hot-swappable and plug-and-play, taking the power from any part of the bus bar to cables leading to the components. Typical products handle currents up to 120A, and consist of open-sided single-pole connectors with typical nickel, silver or gold surface plate, mounted in pairs to provide current loop connection.

It's notable that cloud player Microsoft has also offered hardware through the Open Compute Project, but its Orion project uses AC power, not a bus bar, and has a PDU-like structure. And even this will take a significant investment. Microsoft will be approaching white-label equipment makers, and asking them to meet a specification, but this is not simple.

This is an interesting concept, and Wilson sees the industry moving to a world where "the cabinet is the new blade server."

Servers in 1U or 2U of space share power and network, much like vertically mounted servers in the proprietary blade server systems which emerged in the 1990s. A scheme like that would require a new standard for power distribution.

A couple of years back, IHS analyst Sarah McElroy said it clearly: "The Open Compute Project is something that is peripheral to the rack PDU industry that may prompt more customers to request some level of customization in PDU products which will in turn motivate suppliers to offer more customization."

It's possible that as data center applications consolidate, and technology becomes more standardized and commoditized, more of the world's data centers will evolve into the kind of monocultures which can be built with OCP hardware throughout.

But Wilson says that no matter the level of excitement around OCP hardware, it's actually a long way from any sort of take over: "That would be a very very long slow bandwagon for everyone to jump on."

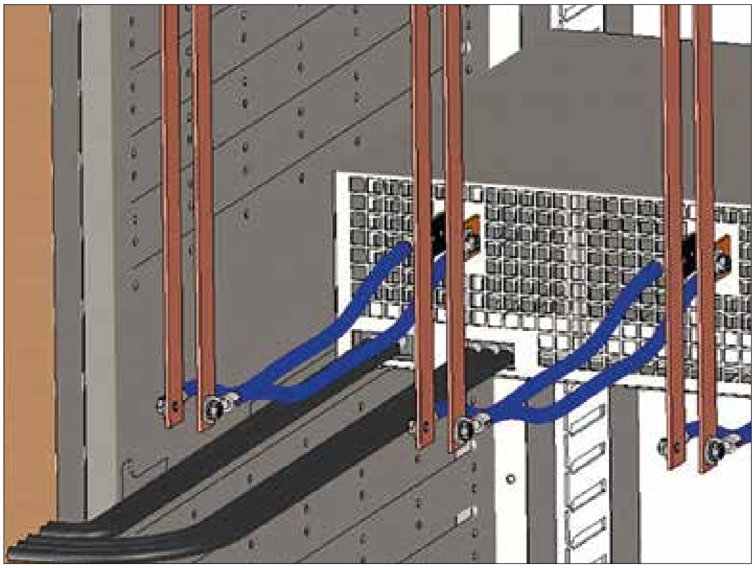

The bus bars in OCP's Open Rack v1

Peter's random factoid

The data center PDU market is predicted to grow by 6.8 percent per year till 2021 (Research & Markets)

Your Operations Team is the Real Winner

AMPPDU

The transformer compartment is convection cooled

Monitor module

Subfeed distribution breakers compartmentalized for arc flash safety

Isolated electronic trip main breaker module compartmentalized for arc flash safety

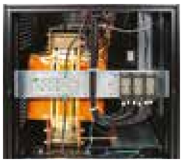

Side mounted distribution compartments on right and/or left side of unit

Transformer taps at 480 and 400 VAC DOE2016 compliant

Meets IEC60349-1 Form 4b Type 7 product that has been tested, approved and UL/CUL Listed

Our PDUs are designed to meet the industry standard for arc flash protection, promoting operator safety for customers, contractors, and service technicians. Our unique compartmentalization technology isolates the input, output, and monitor from the transformer, enhancing safety and helping to eliminate operator errors.

ISOCompartment Design

Our ***ISOCompartment*** design is the safe and efficient way to populate your PDU after it has been powered and installed.

Customers can connect to up to two output breakers in the PDU without entering the *main power cabinet* allowing you to connect new loads in a safe, easy manner.

Separate doors for each ***ISOCompartment***

Two breakers can be left in the off position and locked out until additional power is required.

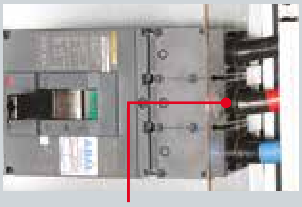

Breakers are pre-wired to the ***ISOCompartment***

3930 Technology Court, Sandston, VA 23150 | anordcpi.com

ANORD CRITICAL POWER INC.

Changing of the Guard: Making the Case for Lithium-Ion Batteries

They're smaller, they're lighter, they last longer... and they're already powering a data center near you. Is it time to switch? By Peter Panfil, Vice President of Global Power, Vertiv

There was a time, not that long ago, when entire buildings were filled with computers performing the types of functions now handled routinely by the phones we hold in the palms of our hand. Technology isn't just becoming more powerful; it's smaller, lighter and, most significantly, it's mobile. We're always connected, part of a real-time global conversation that never stops. The airplane may have shrunk the world, but the smart phone made it pocket-sized.

Powering it all is the lithium-ion battery. Before lithium-ion technologies became viable, energy storage for phones, laptops and tablets was a significant problem. There were crippling limits to just how small these devices could be and how long they could function between charges. Lithium-ion batteries changed everything, providing compact, light, long-lasting energy storage and transforming modern consumer electronics from moderately functional to virtually irreplaceable.

Even as lithium-ion applications became the dominant choice for consumer electronics and increasingly moved into aircraft, automobiles and countless other walks of life, the data center remained the refuge of valve-regulated lead acid (VRLA) batteries. And why not? Data center UPS systems rely on batteries for transitional power. They need to be reliable, and they need to provide just enough power to bridge the gap between utility and generator. VRLA still gets the job done, and is a perfectly fine choice for some data centers.

But there are reasons VRLA batteries are universally regarded as the weak link in the data center power chain. They're big, heavy, they don't hold a charge long, and they need to be monitored regularly and replaced often. The enduring benefit of VRLA batteries is cost, but that gap is shrinking. Today's lithium-ion batteries are about 1.75 times the price of VRLA, on average. That's down significantly in recent years, and there's even a compelling financial case for lithium-ion when you look at total cost of ownership.

Consider this: Lithium-ion batteries store more energy in a smaller space, reducing the battery footprint. They last significantly longer, delaying replacement and reducing replacement costs over time. Most lithium-ion batteries can tolerate higher operating temperatures, reducing cooling requirements and costs. And lithium-ion batteries require less maintenance than traditional VRLA batteries. These are significant and important advantages. Let's take a closer look.

- **Reduced footprint:** Depending on construction and cabinets, lithium-ion batteries can be as much as 70 percent smaller and 60 percent lighter than VRLA, significantly reducing the space required for battery storage. That can reduce construction costs on new builds or increase the amount of usable space in existing facilities. In some cases, lithium-ion batteries can be stored in the row, reducing cable runs.

- **Longer life:** Lithium-ion batteries can last as much as four times as long as VRLA and routinely last two-to-three times as long. Why does that matter? The most significant cost associated with batteries is replacement. In most cases, VRLA batteries would need to be replaced multiple times before the first replacement of a lithium-ion battery.

- **Reduced cooling costs:** There are a number of variables that influence the cooling required for batteries and the associated costs, but most lithium-ion batteries can operate at higher ambient temperatures than VRLA. In some cases, it can reduce battery cooling costs by as much as 70 percent.

The cautious adoption of lithium-ion batteries in the data center isn't surprising. The industry by nature is conservative when it comes to technology change, and echoes of airline and consumer headlines taint any potential discussion of lithium-ion. But the lithium-ion batteries being applied in data centers are a different chemistry – one more akin to that used in the automotive industry – and eminently safe.

Also, VRLA batteries have been a mostly reliable option, and data center managers who stake their careers on reliability and uptime tend to stick with what they know. The problem is, "mostly reliable" isn't exactly accurate and isn't remotely good enough. Understand: A 2013 study from the Ponemon Institute, commissioned by Vertiv, found UPS and battery failure was the leading cause of data center downtime. At some point, the weak link has to be replaced.

As lithium-ion batteries evolve and become more and more viable as an option in the data center, it's becoming increasingly difficult for data center managers to accept the shortcomings of VRLA. New chemistries and construction practices have improved battery safety, reduced capital costs and rendered some of the arguments against lithium-ion obsolete.

Bottom line: Today's lithium-ion batteries are safe, reliable alternatives to VRLA with a compelling TCO case.

VERTIV™

Contact Details

Phone: *614-888-0246*
Email: *Contact@VertivCo.com*

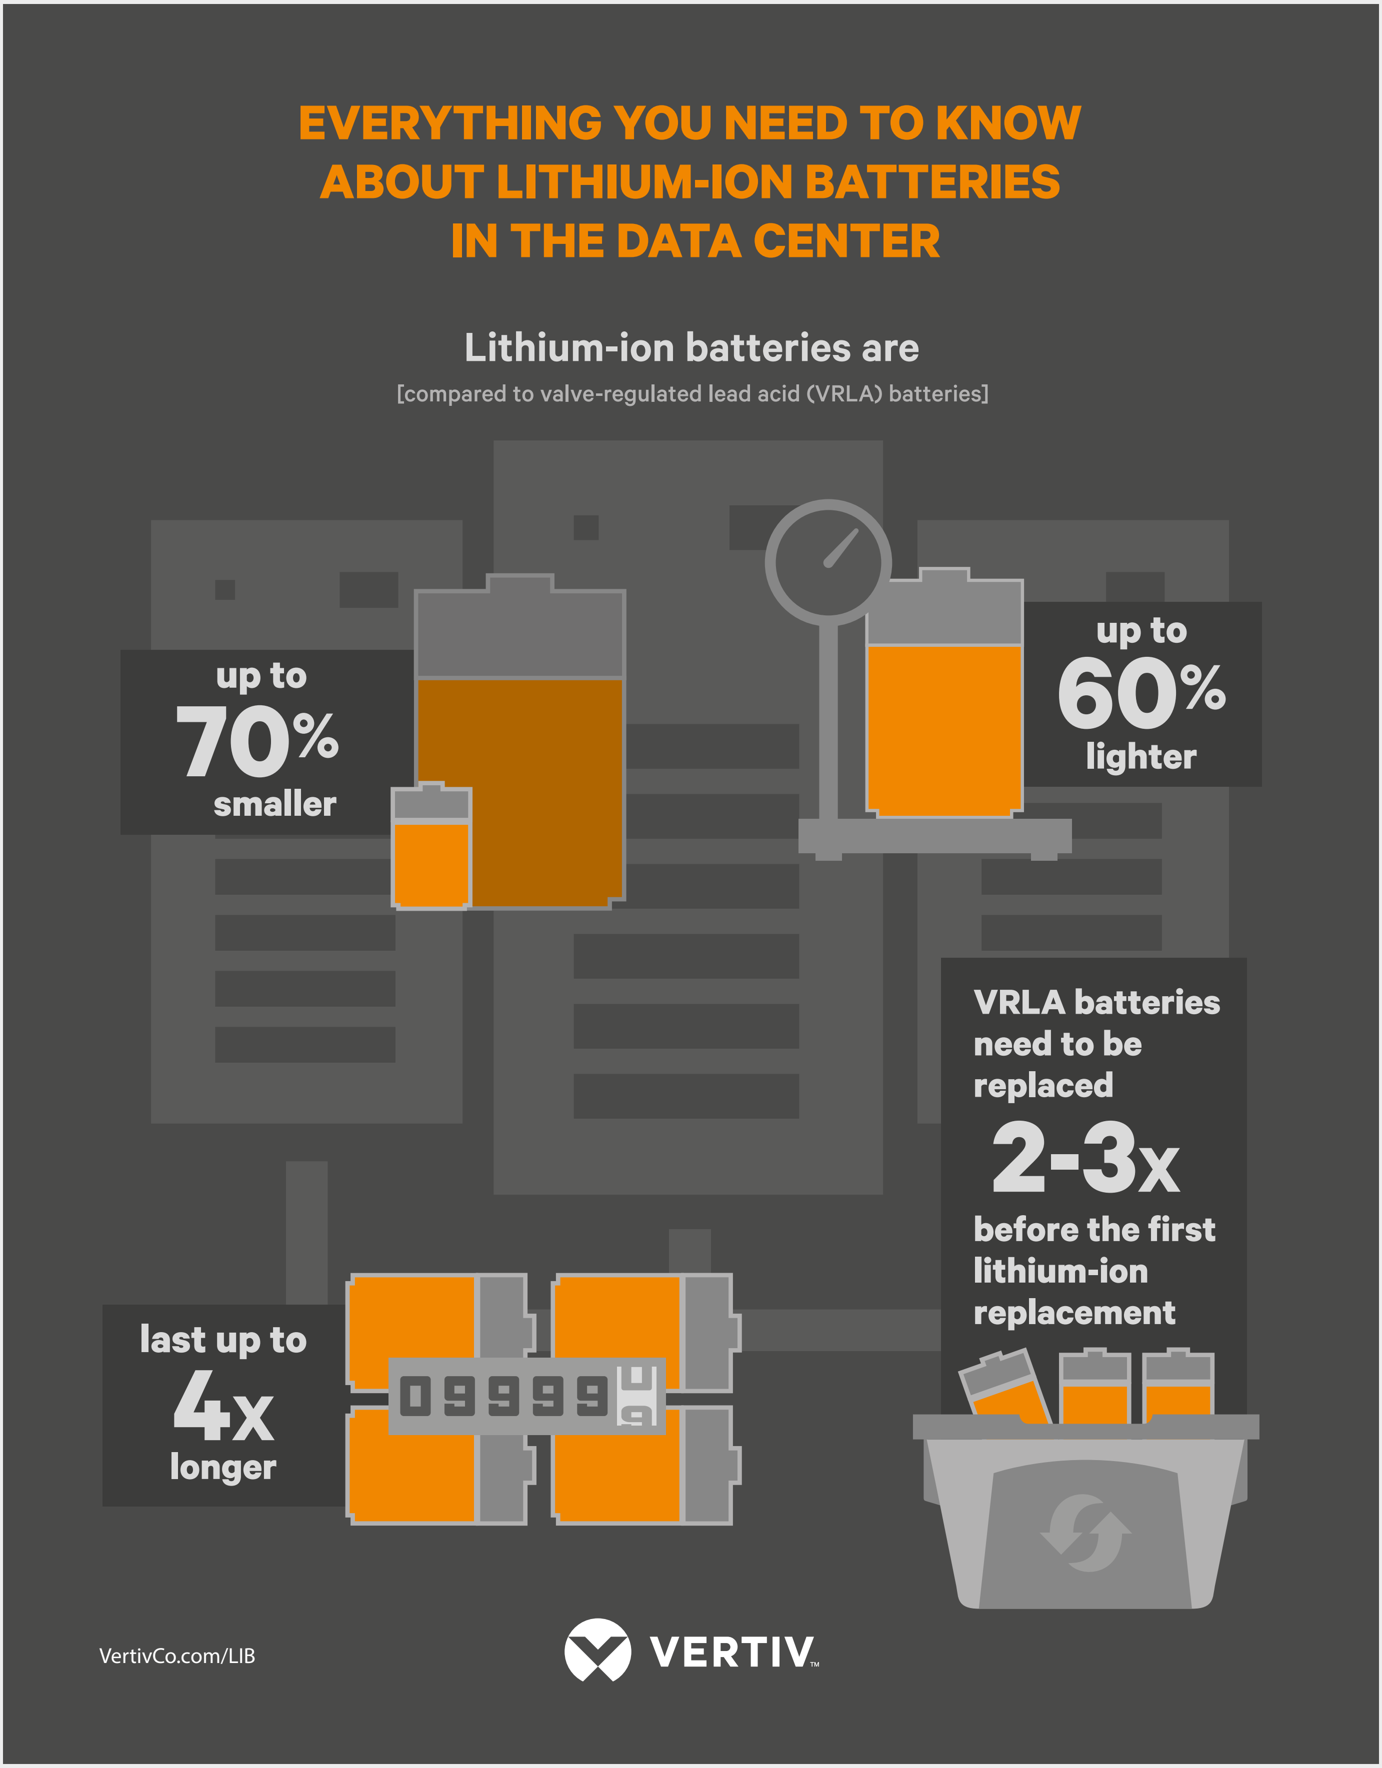
EVERYTHING YOU NEED TO KNOW
ABOUT LITHIUM-ION BATTERIES
IN THE DATA CENTER
Lithium-ion batteries are
[compared to valve-regulated lead acid (VRLA) batteries]
up to 70% smaller
up to 60% lighter
VRLA batteries need to be replaced 2-3x before the first lithium-ion replacement
last up to 4x longer
VertivCo.com/LIB
VERTIV

Power chips, but not as we know them...

Max Smolaks welcomes a new material which will replace silicon in the power chain

Max Smolaks
News Editor

For the past 35 years, most power supplies have relied on power MOSFETs (metal oxide semiconductor field effect transistors) – voltage-controlled devices made of silicon that are used to switch and condition electricity.

These little black squares have done a great job, scaling their performance with Moore's Law and making their way into all sorts of data center equipment, from PSUs to routers, switches and servers.

But the entire power MOSFET family is about to become history, killed off by a class of devices that will be smaller, more efficient and cheaper too – at least in the long run. Enter gallium nitride (GaN), the wonderful semiconductor that is going to cut your electricity bills.

Thanks to a precedent set by Alexander Graham Bell, intermediate voltage comes into a data center at 48V and must go through multiple stages of power conversion before it reaches components on the board, losing a portion of its useful energy at every stage.

"Silicon wasn't fast enough to get from 48V all the way to 1V," Dr. Alex Lidow, chief executive of Efficient Power Systems (EPC) and one of the inventors of the original power MOSFET told *DCD*.

"So what we [as an industry] did was create a whole bunch of very expensive power supplies that get you from 48V to 12V, and another set of power supplies that get you from 12V to 1V. And with gallium nitride, since it's so damn fast, you can get rid of that whole intermediate bus and go directly from 48V to 1V."

In 1999, Lidow became the chief executive of International Rectifier – the world's oldest independent power semiconductor company, established by his father Leon. In 2014 International Rectifier was sold to Infineon for $3 billion, liberating Lidow to focus on his other venture, built on the belief that gallium nitride would change the world of electronics.

EPC has been manufacturing GaN chips since 2000 but the initial production costs were prohibitively high. Early applications included LIDAR lasers, wireless power transmitters and the 'colonoscopy pill' - the first ingestible imaging capsule that uses low-dose X-rays for cancer screening. As the costs went down, EPC turned its attention to data centers.

Fig.1. EPC's GaN power MOSFET structure employs an AlN isolation layer between the silicon substrate and the GaN (below)

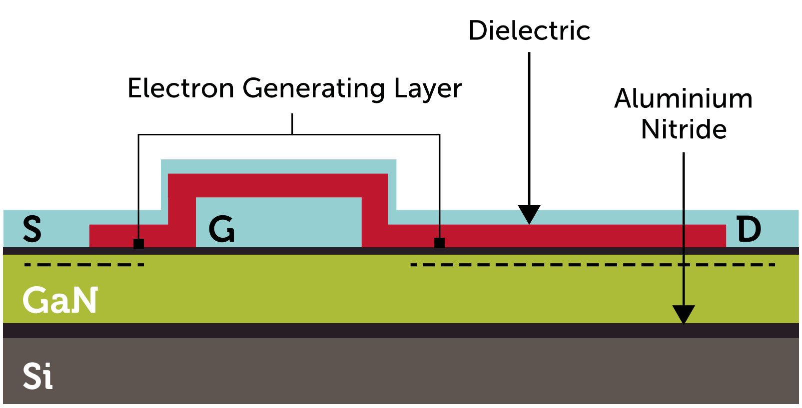

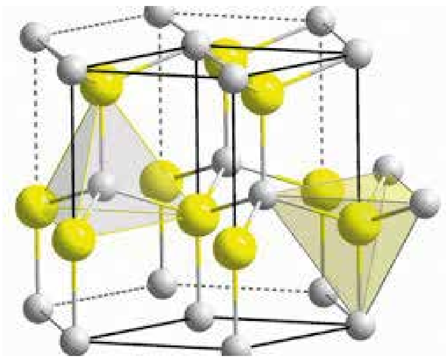

Fig.2. Gallium nitride has a Wurtzite crystal structure, and a wide band gap of 3.4eV, useful in high frequency devices

"Now it's becoming a significant market for us. It has taken many years to convince data center people that GaN is reliable enough – and to get our costs down to the point where it's cost-effective for them to take the plunge," Lidow said. "Maybe even more telling, Texas Instruments uses our product inside their products for data centers."

48V power distribution throughout the rack is a hot topic, being investigated by major data center operators like Google, Facebook, AWS and Microsoft. Lidow told us most hyperscale companies are already testing GaN chips in their equipment - without naming any names.

According to EPC, transmission and power distribution losses mean that today, keeping a 580W machine switched on actually requires 860W of power to be produced. With GaN chips implemented every step of the way, we would need just 770W at the source.

It is not just power efficiency: GaN chips are also much smaller than their silicon counterparts, and produce less heat. Traditional power distribution components, together with the attendant heat sinks and fans, occupy valuable space on a server board that would be better spent accommodating additional CPU cores or RAM modules.

"It used to be that the server was like a mouth, ears and a brain – a lot came in through the ears, went out the mouth, but the brain didn't have to do very much. Now with artificial intelligence and cloud computing the brain is really cranking, and there's not that much more coming in through the ears or going out the mouth," Lidow said.

"What this means is a server has to communicate internally a whole lot more, and that puts a different metric on performance: there's a big push to condense the boards and pack the servers much tighter, so that you can get this thinking process going. That is limited by the heat and power density of these systems, so you need more efficient power conversion – that's one of your largest sources of heat."

The manufacturing process for GaN chips involves growing a thin layer of gallium nitride on a standard piece of silicon, but the chips themselves are much smaller, therefore you get more chips for every manufacturing batch. During manufacturing, GaN is encapsulated in silicon, so it doesn't require packaging - or additional protection.

"Moore's Law kinda ran out of gas in power distribution before it ran out in digital," Lidow said. "But with gallium nitride, our first chips were five to ten times better than the theoretical performance of silicon. We've been doubling that performance every few years, and we're about to double it again."

And there's still room for growth: "Even with all that doubling in performance, we are still going to be around 300x away from theoretical performance of GaN."

Next up for GaN chips, it's the long slog up the supply chain. Texas Instruments uses EPC's chips in products like the LMG52000 half bridge power stage, which are sold on to firms making power supplies, which are built into servers, storage and networking equipment. Hyperscalers could fast-track this, by adding components to their own hardware built to exact specifications.

> *A 580W system needs 860W of raw power. GaN could cut this to 770W*

Another company using gallium nitride for power distribution is GaN Systems, established in 2008 to capitalize on research carried out by now-defunct Nortel.

"The entire industry recognizes that GaN is the future, it's much better than silicon, it's not an 'if', it's a 'when,'" Paul Wiener, VP for Strategic Marketing at GaN Systems told *DCD*.

Wiener thinks that 'when' can be brought much sooner if C-level executives at giant data center customers like Google, Facebook, Amazon and Microsoft shift their focus.

"The incremental purchase price is drowned out by the savings in opex and additional revenue," he said. "Servers generate revenue - bits per dollar. The more data that those servers can process, and more servers can be put into a rack, the more revenue that rack generates."

According to calculations published by GaN Systems, switching to gallium nitride can increase server space capacity by 14 percent. And that alone, without taking higher efficiency of chips into account, can boost the 'incremental revenue potential per rack' over three years by $31,686 - with one-time capital investment already subtracted.

Early adopters can buy equipment powered by GaN chips from a number of lesser known vendors: they will see considerable efficiency gains, but at a higher initial cost. GaN Systems expects its own chips will reach price parity with silicon-based power MOSFETs in 2019 or 2020 - making the latter all but obsolete.

Chasing lithium nirvana

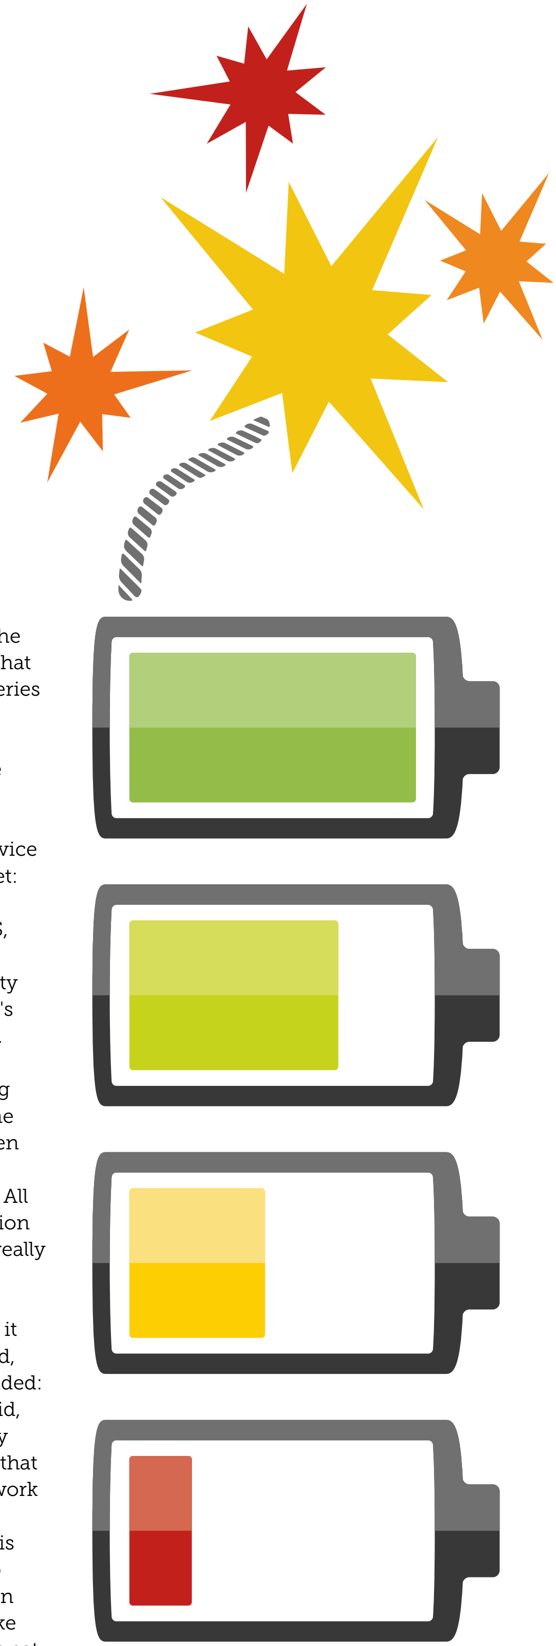

Lithium-ion batteries could save money in data center backup, says *Sebastian Moss.* The trouble is, they have to live down the exploding Galaxy Note 7

Sebastian Moss
Reporter

The coming lithium-ion UPS revolution has been prophesied for some time, with lead acid batteries long seen as a technology of the past. On paper, it's easy to see why people have thought this - Li-ion is lighter, it's smaller, and it lasts longer. Surely, it is the superior technology, one might think.

And yet, here we are in 2017 and the only viable competitor to valve-regulated lead acid (VRLA) seems to be the flywheel. Why is this?

One of the major problems some have had with Li-ion is its tendency to explode. "Just like with terrorists, there are too many bad stories of Li-ion in the news, people are intimidated," Schneider Electric's Simon Zhang told *DCD*.

"People see explosions everywhere, from hoverboards, to Boeing, to cellphones."

Perhaps most prevalent on the public's mind is what happened with Samsung's Galaxy Note 7, which had to be recalled after units kept catching fire.

"For this reason people are cautious, especially for mission critical facility operators and managers. They're a very cautious, very conservative group of people," Zhang said.

"They don't want to risk their facility with this relatively new technology, even though it's been there for 30 years."

Peter Stevenson, senior technical coordinator at GS Yuasa Battery agreed: "If you lose a data center, the costs are so much more than anything you're going to save on fitting a different battery, so I can see why the reticence is there."

Indeed, the idea of filling one's data center with a host of giant Galaxy Note 7s could seem rather off-putting, but "that's why it's mandatory to have a battery management system (BMS)," Zhang said. "If anything goes wrong, or there are any safety concerns, it will shut off the battery, take that string of batteries offline."

This safety feature differentiates it from the exploding device in your pocket: "Because we have the BMS, we have the whole visibility of the battery's performance. You can see the remaining capacity of the batteries, when it needs to be replaced, etc. All this information will become really valuable."

It also differentiates it from lead acid, Stevenson added: "With lead acid, really the only way to know that it's going to work and perform as necessary is to actually do a discharge on it - it's a bit like Schrödinger's cat. But if you do that every week, then it's not going to last so long.

"Li-ion will soak up a cycle a day for ten years quite happily, so you can see how things are changing day by day, so really

Li-ion is the ultimate security battery," said Zhang.

Unfortunately, if done poorly, a BMS can bring its own risks, particularly due to incompatibility, said Zhang: "I heard that the Boeing incident was caused by the battery vendor and the BMS vendor being different."

This difference between vendors is an issue still plaguing the nascent industry, Zhang continued: "At this stage it's a pretty fragmented market, the batteries are not standardized, even the voltages are different, the chemistry, the form factor, technical specifications, everything is different."

"That's why you see so many incidents, all due to the fact that the safety standards are different from manufacturer to manufacturer."

Stevenson concurred, adding that "there are many alternatives, but they all use the same principle of shuttling lithium-ions backwards and forwards."

There's a mass of different materials that can be used, and they're developing all the time, and they have some quite different characteristics which make them suitable for different applications. If you only make one type you tend to push that for everything."

There is a debate in the industry over which type of Li-ion battery is best for which scenario, but Zhang is hopeful that the sector will organize itself: "Let the free market help us to filter through and consolidate to two or three vendors."

But if you do trust your vendor and the solution that they offer, it is clear that Li-ion does bring some genuine benefits to the data center. "The general rule of thumb is that we save two thirds of the footprint and weight," Zhang said.

The battery lasts longer, too, on average chugging along for 10-15 years, rather than three to six years for lead acid systems. "Another good thing is that it doesn't collapse like lead acid. With lead acid, towards the end of life, you can get some sudden failures," Stevenson said. "With Li-ion you get a gradual deterioration that can be measured quite simply."

Li-ion can also bring cost savings as it can run efficiently at higher temperatures. While VRLAs perform optimally at 20 degrees, Li-ion can operate in temperatures closer to that of an average ambient environment.

"So you've got big savings on HVAC systems," Stevenson said.

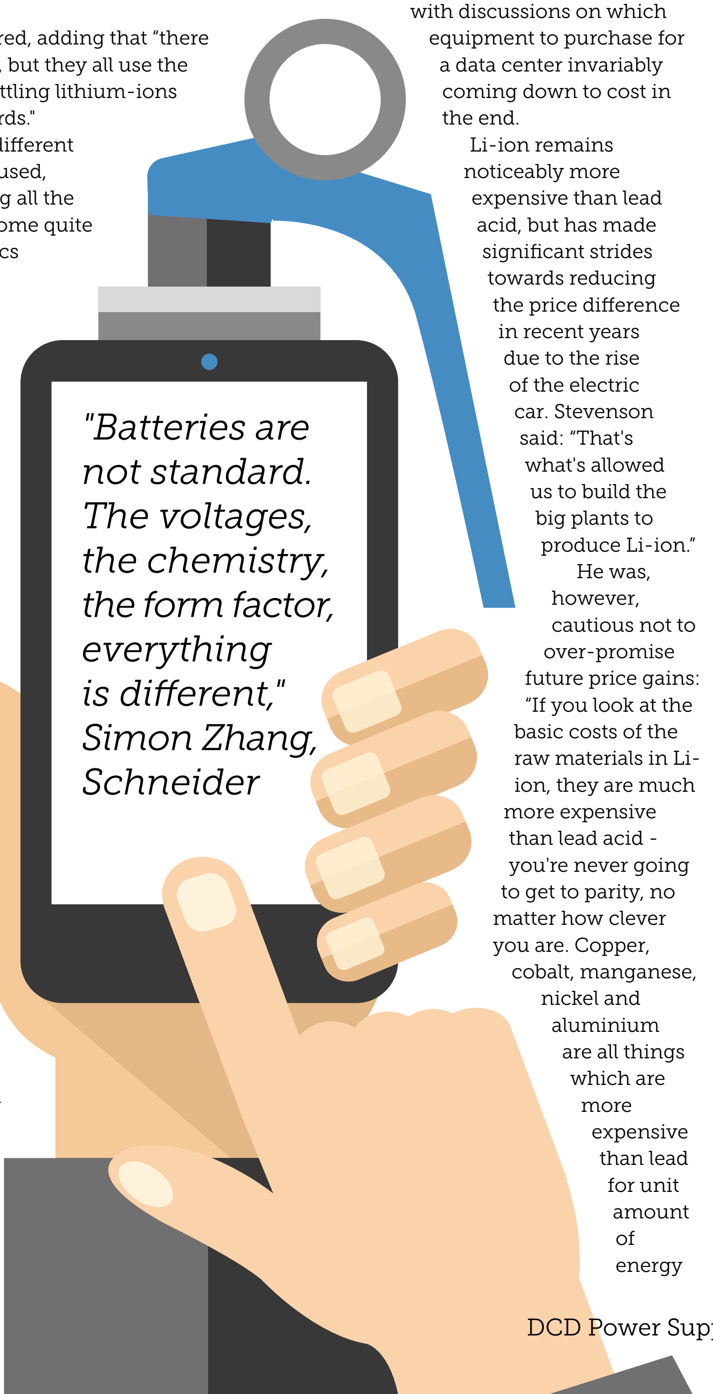

Savings are important, with discussions on which equipment to purchase for a data center invariably coming down to cost in the end.

Li-ion remains noticeably more expensive than lead acid, but has made significant strides towards reducing the price difference in recent years due to the rise of the electric car. Stevenson said: "That's what's allowed us to build the big plants to produce Li-ion."

He was, however, cautious not to over-promise future price gains: "If you look at the basic costs of the raw materials in Li-ion, they are much more expensive than lead acid - you're never going to get to parity, no matter how clever you are. Copper, cobalt, manganese, nickel and aluminium are all things which are more expensive than lead for unit amount of energy stored."

But people should not look at just the ticket price of the battery, both Stevenson and Zhang agreed. Data center operators should look at the long term savings that a Li-ion battery can offer. Unfortunately, as of yet, few make this comparison.

"You can make a lower Total Cost of Ownership (TCO) case, but they still buy on capex, or up-front cost," Zhang said. "If it's two times the cost, it's hard to justify, because normally the buyer is different from the people who run the facility."

Stevenson said that "some of the telecoms companies have actually put their money where their mouth is and looked at 15 year whole life costs, but very few people are willing to look more than four or five years. They want the payback in that time."

Time is what will prove the true arbiter on whether Li-ion will manage to make a mark in the data center, as operators wait to see how the market unfolds, each hoping for someone else to take the plunge and test out Li-ion at a large scale. If that happens successfully without facing any dramatic issues, it could pave the way for much wider adoption.

"We've been trying to get the message across, but so far I don't think it's been swallowed," Stevenson said. "It needs one or two of these data centers to go big time, and people will understand it and see the benefits."

"It's on a cusp," he continued. "We could be on that cusp for many years, or it could change very quickly. It really depends on if people can be convinced."

Zhang said: "Education is the first thing we need to do. We see some early adopters, we've got a few. If they install them, that will help convince later adopters that this is a safe technology to use." He added optimistically: "We feel like we're at the turning point of switching to this new technology."

But for those hoping for a rapid transition of power storage, Stevenson offered a sobering anecdote.

His company had recently been chosen to supply Li-ion batteries for the International Space Station - a hugely valuable contract - but getting the world's space agencies to feel comfortable putting a device sometimes associated with explosions on man's single most expensive construction in history was no easy task.

"We've been putting Li-ion into space for 20 years, but they certainly didn't put them in the manned space stations until just recently," he said. "It's taken that long for people to change for manned space flight, and you work up to it through less critical situations first.

"You can't force things, you have to wait for people to make up their mind."

Carborundum may grind out eco-mode UPS

Illustration: studionippoldt.com

Eco-mode saves power but increases the risks associated with a UPS. *Ian Bitterlin* says silicon carbide could give the best of both worlds

Power losses in static UPSs have gradually decreased over the past 20 years. Early thyristor based on-line systems (referred to as double-conversion or IEC 'VFI') with transformers at the input and output had full-load operating efficiency of 83-85 percent. Now to transistor (IGBT) based line-interactive (VI) transformer-less machines achieve 97.5-98 percent.

Energy efficiency has increased by 15 percent while cooling demands have fallen. Reliability has increased, with a module MTBF rising from less than 25,000h to more than 150,000h. Output voltage waveform distortion has gone down from five percent to one percent. Noise is down from 95dBA to 70dBA, and the physical footprint is down by a staggering 90 percent.

Even double-conversion (VFI) has reached 96.8 percent efficiency, and the cost per kW capacity has fallen to its lowest level ever.

This is good for the purchaser, but the only way to make a profit from UPS is to provide after-sales-services. By 2008, where was there left to go?

In Europe, all UPSs had already become transformer-less, thanks to Europe's four-wire distribution, while in North America, transformer-less UPSs are still a minor novelty and often regarded as somewhat exotic. At the same time, vendors like APC adopted line-interactive topology (IEC 'VI') which saved energy, albeit without any frequency protection – not technically 'on-line', although advertised as such, but working well enough in stable grids.

'Eco-mode', introduced by Invertomatic in Switzerland in the 1990s but dropped due to lack of sales, was resurrected (along with 'modular' UPS, which overcame partial load problems endemic in most data centers). The principle of eco-mode is simple: when the utility is stable, the UPS switches itself into bypass mode and the losses reduce, especially in transformer-less designs. The rectifier still floats the battery (needing far lower power than a flywheel) but the inverter is throttled right back and, in the best designs, the cooling fans are dropped off.

The automatic bypass (a thyristor switch) keeps the load on the utility until the utility shows the first sign of deviation – at which point the static switch transfers the load back to the inverter, all in under 4ms and within the (rather outdated) ITIC/CEBMA PQ Curve. The UPS then monitors the utility for stability and, after a period, usually one hour, switches the load back to bypass. The advantages are clear; 99 percent efficiency for more than 95 percent of the year on stable grids, with the bonus of excellent low-load efficiency.

> *Who will need to worry when you can get protection with one percent loss?*

Some unscrupulous salesmen mention 'low-power state' for the inverter but, make no mistake, the UPS is in bypass with no power quality improvement and the critical load fed by 'raw mains.'

Now, there are some 'advanced eco-modes' which operate at 2ms instead of 4ms, and some that monitor the load distortion and make decisions about the grid, but the basic concept remains – if the utility is stable you save energy.

Rewards usually come with risk and eco-mode is no different. Every time the utility deviates, the load is switched – the very opposite of the protection offered by 'double-conversion.' This switching represents a risk to the load, which may be small but the user must balance it with a return - which can be high enough to cover the entire UPS cost in less than two years.

As energy costs rise and the concept is proven, eco-mode is being accepted. Energy effectiveness is not always the most important metric that users aspire to, but there are even a few high reliability dual-bus facilities that are hedging bets by enabling eco-mode in one bus and running VFI in the other, alternating each week.

The risk, whether real or perceived, will remain and limit the adoption of eco-mode but a development emerging from Japan could negate any advantage of eco-mode.

Transistors are currently manufactured with layers of doped silicon. The best to date, for UPS, are of the Insulated Gate Bipolar (IGBT) type and have become increasingly powerful and reliable.

One drawback is that the faster you switch them (to achieve more precision), then the higher the losses.

This is what mainly contributes to the upper limit of 96.8 percent module efficiency.

However, a change from silicon to silicon carbide (better known as carborundum or occurring in nature as the extremely rare mineral moissanite) can increase UPS module efficiency to 99 percent in double-conversion.

Synthetic silicon carbide powder has been mass-produced since 1893 for use as an abrasive, for instance in silicon carbide paper for finishing metals.

Silicon Carbide IGBTs will initially cost more but the energy saving will rapidly be recovered – and all without switching the critical load to the raw utility and increasing risks of transfer.

Hence, silicon carbide will spell the end of worrying about the enablement of eco-mode and possibly even kill off line-interactive (VI) UPS. Who will need to worry when you can get total protection of voltage and frequency protection with less than 1 percent loss?