

Designing a “superscanner” for high-throughput battery QC: Technical overview of the world’s fastest industrial CT scanner



Authors:

Peter Attia, cofounder and CTO, Glimpse (peter@glimp.se)

Emil Espes, Head of Segment, Excillum (emil.espes@excillum.com)

Christian Gück, Head of Sales, VCbattery (c.gueck@visiconsult.de)

Version: V1 (2025-09-29)

Table of Contents

Summary.....	3
Introduction.....	4
CT scanning primer.....	5
Superscanner design objectives.....	6
The Excillum MetalJet X-ray source.....	8
X-ray sources and the thermal management challenge.....	8
The Excillum MetalJet.....	10
The VCB2 “superscanner”.....	14
Imaging chain: X-ray detector.....	14
Imaging chain: CT acquisition modes.....	15
Imaging chain: Mechanical and thermal control.....	17
Imaging chain: Reconstruction.....	17
Sample scans.....	17
Part handling.....	18
Production considerations: Reliability, serviceability, and automation.....	20
The Glimpse Portal.....	22
Conclusion.....	28
References.....	30

Summary

Computed tomography (CT) scanning is invaluable for battery quality control, but its speed has historically been insufficient to match the unforgiving throughput of battery production. Fortunately, recent breakthroughs across the CT technology stack have unlocked a step change in CT speed. In this white paper, we discuss how [Excillum](#), [VCbattery](#), and [Glimpse](#) designed and built the “superscanner”, capable of high-resolution, full-cell cylindrical cell scans in 5–10 seconds. Specifically, we’ll discuss how Excillum’s ultrabright X-ray source, VCbattery’s system integration and reconstruction, and Glimpse’s battery quality management solution dramatically accelerate time-to-insights on battery quality for cell producers and OEMs alike. This white paper may be of interest to battery engineers, quality engineers, and/or inspection industry professionals interested in understanding the design and implementation of the world’s fastest industrial CT scanner.

Introduction

The battery industry faces three key challenges as it scales. Briefly, battery safety is high-severity yet often caused by subtle cell defects; battery reliability is threatened by the vulnerability of packs to single-cell failure; and battery manufacturability is impeded by the stringent requirements for purity, tolerance, yield, throughput, and cost. All three problems are closely related to battery quality; thus, enabling battery quality at scale is a major challenge.¹

While battery quality is a challenging problem requiring a multifaceted solution, a key lever to improving battery quality in production is advanced cell inspection. Among the various techniques available for battery cell inspection, [computed tomography \(CT\)](#) is often considered the gold standard as it provides 3D, nondestructive², full-cell evaluation with high spatial resolution.¹ Historically, however, CT has been limited to R&D and failure analysis applications due to its slow speed—specifically, both its slow scan time (often ~hours) and its slow analysis time (also often ~hours) per cell. While recent advances have driven down the scan time to ~minutes per cell³, a step change in speed is required to make CT a viable tool for cell production.

Over the past two years, three companies—[Excillum](#), [VCbattery](#) (a division of [VisiConsult](#)), and [Glimpse](#)—have co-developed what we call the “superscanner”, the world’s fastest industrial CT scanner. Formally known as the VCB2, the superscanner combines several state-of-the-art advancements across the CT tech stack including [Excillum’s MetalJet X-ray source](#)⁴, [photon-counting detectors](#)⁵, VisiConsult’s system integration and part handling automation expertise, high-speed reconstruction, and Glimpse’s data processing, management, [visualization](#), and automated inspection solution. As a result, this system can achieve high-resolution, full-cell CT scans of cylindrical cells in 5–10 seconds, which translates to ~10,000 scans per day.


In this white paper, we provide a technical overview of the “superscanner”. Specifically, we detail the objectives, requirements, and design of this groundbreaking CT system and its state-of-the-art components. This white paper is recommended for battery engineers, quality engineers, and/or inspection industry professionals interested in understanding the design and implementation of the world’s fastest industrial CT scanner.

CT scanning primer

If you're new to CT scanning, here's a one-page primer.

[Computed tomography \(CT\)](#) is a non-invasive imaging technique that uses X-rays to create detailed cross-sectional images of the inside of an object. You might know CT by its old name, "computed axial tomography", or "CAT" scanning. In short, X-ray CT is 3D X-ray imaging. You've likely encountered CT imaging if you've ever had a medical CT scan for an injury or illness, or if you've had your luggage inspected during an airport security screening.

In X-ray imaging, X-rays are first generated by an [X-ray source](#). The X-rays pass through the object of interest and are detected by an [X-ray detector](#). The object rotates on a stage during scanning, producing a sequence of hundreds or thousands of 2D X-ray images. These 2D X-ray images (termed **radiographs** or **projections**) are then [reconstructed](#) into a 3D volume that approximates the original object. A schematic of the basic geometry of a CT scanner is displayed below (taken from Glimpse's [FAQ page](#)):



See the below references for more details:

- https://en.wikipedia.org/wiki/Industrial_computed_tomography
- <https://rigaku.com/products/imaging-ndt/x-ray-ct/learning/x-ray-computed-tomography-brief-introduction>

CT image quality is a complex topic in its own right. Read [here](#) to learn more.

Superscanner design objectives

The key design objective for the superscanner is speed. An overly simplified interpretation of this objective is simply “scan time”. Obviously, scan time is a critical attribute of the superscanner. However, we consider “time-to-insights” to be the best measure of overall system speed, since fast scans are useless without the ability to extract insights from the scan data at a commensurate rate. The key factors determining time-to-insights include:

1. **Scan time**, with sufficient image quality to detect the features and defects of interest
2. **Cycle time**, or the sum of scan time and the “dead time” due to part handling per scan
3. **Reconstruction & image processing time**
4. **Algorithm inference time** for features and defects

Each of these factors must keep pace to maximize the system’s utility and to ultimately minimize the time to draw conclusions from CT data.

Since defects and features of interest can occur anywhere in the cell, this system is primarily designed for full cell scanning, as opposed to scanning a small region of interest such as the overhang region.

The superscanner faces additional design considerations as a system intended for production environments. First, this system must be designed for extremely high reliability and serviceability. Additionally, this scanner must be able to seamlessly integrate with automated cell handling equipment. Furthermore, the footprint of this system must be as small as possible, as factory floorspace is often at a premium; this limitation alone, let alone system cost, eliminates the possibility of using multiple “conventional” (i.e., current-generation) CT scanners.

We now discuss three aspects that are out-of-scope for this system:

1. This system is only designed to scan cylindrical cells, specifically 18650 to 46XX cells. The maximum supported cell height is 120mm, although this limit can be adjusted if required. Cylindrical cells are a natural choice for the first superscanner system as they are small, have circular symmetry (i.e., well-suited for CT), and are relatively standardized within the battery industry. That said, this design concept could translate to pouch and prismatic form factors as well—not to mention non-battery applications.
2. This system is designed for ultra-high-throughput cell-level scanning, with voxel sizes on the order of 25 μm . Even with longer scan times, this system is only capable of scanning within a narrow voxel size regime of 20–30 μm . In technical terms, this system is a micro-CT system, which means that the desired resolution is on the order of 5–100 μm as opposed to <1 μm (“Nano-CT”) or >100 μm (“Meso-CT”). Specifically, the prefix refers to the focal spot size range of the X-ray source (i.e., nanofocus, microfocus, mesofocus). Read more [here](#) and [here](#).
3. While the superscanner can scan 6–8 cells per minute (~10,000 cells per day), a large cylindrical battery production facility produces 100s to 1000s of cells per minute (millions per day).⁶ Thus, this system’s throughput is still far below that required for 100% inspection for these large cylindrical cell producers, meaning its utility as a simple pass/fail tool is limited. Instead, this system is intended to operate in conjunction with other production datastreams, such as a manufacturing execution system (MES), to intentionally sample cells across production and root cause any issues that may arise. We’ll return to this point later in this white paper.

We’ll begin by discussing the innovation at the heart of the superscanner: Excillum’s MetalJet X-ray source.

The Excillum MetalJet X-ray source

Before we discuss the MetalJet specifically, we'll review the general principles of X-ray sources and the key challenge in their design.

For a high-level overview, we recommend reading [Glimpse's blog post](#) describing three X-ray source designs. You can also deep dive this topic with the following references:


- https://en.wikipedia.org/wiki/X-ray_tube
- <https://www.excillum.com/history-of-x-rays-early-years/>
- <https://www.excillum.com/history-of-x-rays-x-ray-tubes/>

X-ray sources and the thermal management challenge

At a very high level, an X-ray source has a cathode and an anode just like a battery; unlike a battery, the cathode and anode are separated by air in a vacuum environment. During operation, current passes from the cathode to the anode. X-rays are generated when electrons hit the "target" component of the anode assembly. Typically, 99% of the electron energy converts into heat, and only 1% of the energy will convert into X-rays. Thus, the primary challenge in X-ray source design is maximizing the thermal loading of the target.

A common figure of merit for X-ray sources is electron beam power for a given spot size. Defining both terms:

- [Electron] beam power, measured in watts (W), is the product of beam voltage and beam current. To a zeroth-order approximation, higher electron beam power = more photons (the reality is more complex, but out of scope for this white paper).
- Spot size refers to the size of the area on the source's anode where the electron beam hits. This area is where the X-rays are generated. The smaller the spot size is, the more that the X-ray source approximates a [point source](#)—which directly influences image sharpness. In technical terms, a smaller spot size enhances image resolution by reducing [geometric unsharpness](#), making it ideal for detailed imaging tasks. The larger the spot size, the larger the [penumbra](#) in the image. The following diagram shows how focal spot size affects geometric unsharpness:



Left: Imaging with a point source, or a focal spot size = 0. **Right:** Imaging with a non-point source, or a focal spot size > 0. As the spot size increases, the edges blur due to the [penumbra effect](#).

Ideally, we want a source with high power (more photons) and a small spot size (higher sharpness). Unfortunately, power and spot size are closely linked. The limitation here is practical: If too much electron energy is focused on a single spot of the anode, the (tungsten) metal target will melt! The first half of [this video](#) nicely illustrates this effect.



As a result, the rule of thumb for conventional X-ray sources (“static targets”, where the anode is simply a solid piece of metal) is that the spot size increases by 1 μm whenever the power increases by 1 W. In other words, we must trade off power (\sim X-ray brightness) and spot size (image sharpness) in practice; this tradeoff is one reason why high-resolution scans are generally slower than low-resolution scans (the brightness must be lower since the sharpness requirements are stricter for high-resolution scans).

One design to improve this brightness–sharpness tradeoff is to use a “rotating anode target”. A rotating anode target more efficiently removes heat by, well, rotating the anode target. This way, the focal spot is continuously regenerated; the target is actively cooled, and thus a small spot size can support higher power without running the risk of melting the source’s target. The result of this innovative approach to thermal management is that instead of a 1:1 relationship between power (W) and spot size (μm), a typical rotating target microfocus

source has a 3:1 relationship. The end result of higher brightness at the same spot size is faster scanning at iso-image quality.

The Excillum MetalJet

Excillum's [MetalJet](#) technology takes X-ray source thermal management to an entirely new level. The MetalJet's key innovation is to use a liquid metal jet as the target. This liquid metal jet is continuously regenerated (at nearly 100 meters per second!) and can thus remove heat at an extraordinary rate. [This video](#) from Excillum perhaps best illustrates this innovation. The [MetalJet F](#), [launched in March 2025](#), is the latest generation of this technology.



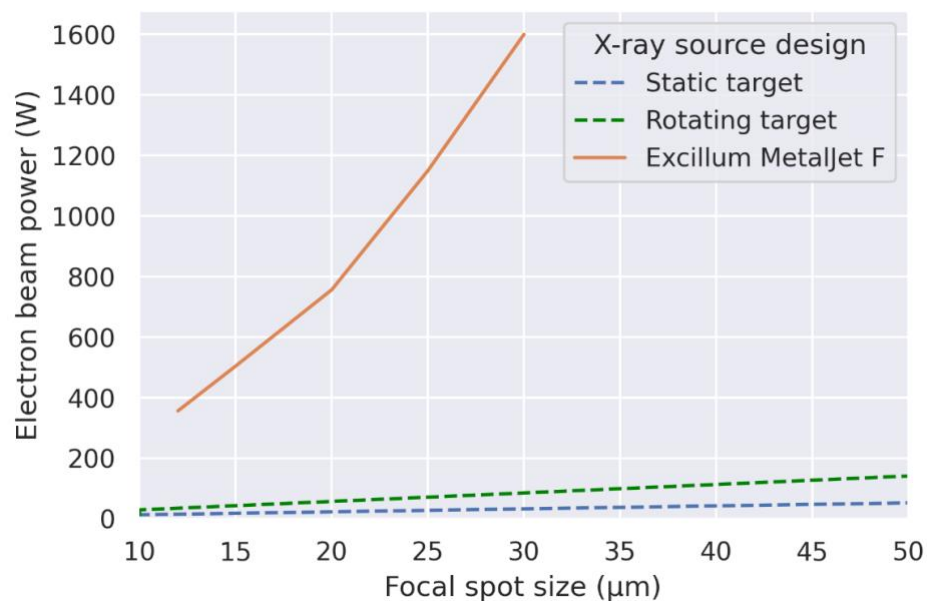
Left: A conventional static-target X-ray source. The target begins to melt at higher electron beam currents. **Right:** A liquid-metal-jet X-ray source. The target is capable of highly efficient heat removal.



The [Excillum MetalJet F1016](#).

The benefits of this approach are massive: this X-ray source yields up to **27x** higher brightness than a conventional X-ray source and $\sim 9x$ higher brightness than a typical rotating anode source. Again, to a first order approximation (namely, ignoring the critical role of the X-ray detector), scan time is inversely proportional to X-ray source brightness.


The following plot compares the three X-ray source designs we've discussed so far in terms of electron beam power vs. focal spot size. An ideal X-ray source will have high electron beam power at a small focal spot size. This plot clearly illustrates how the MetalJet is truly a breakthrough X-ray source.



Electron beam power vs. spot size for three different X-ray source designs: static target sources, rotating target sources, and the [Excillum MetalJet F1016](#). In this plot, high electron beam power and small focal spot size is ideal. The MetalJet truly represents a breakthrough over existing X-ray source technologies.

MetalJet sources have another major advantage over both static and rotating target sources: they produce two X-ray beams for only one electron beam. X-rays radiate in all directions from the focal spot where the electron beam strikes the target. Thanks to the MetalJet's symmetrical design with two opposing windows, both sides of the same focal spot can be accessed, providing two full-strength beams instead of one. When paired with two detectors, this dual-beam design translates to an additional 2x throughput increase.

While the MetalJet has only been recently introduced into CT scanning applications, Excillum itself was [founded in 2007](#) and has deployed hundreds of MetalJet sources across diverse X-ray applications including diffraction, scattering, fluorescence, and more. Systems incorporating prior MetalJet designs have been deployed in both R&D and production environments. The MetalJet F has incorporated field feedback from these systems and has been designed for ultrahigh reliability. One of the most impressive reliability-oriented innovations introduced in this generation of the MetalJet is electromagnetic pumping, which enables pumping of the liquid metal with no moving parts. Additionally, the system is designed for 24/7 operation with a continuous duty cycle (i.e., the liquid metal is continuously pumped). Furthermore, because the liquid metal target is continuously regenerated, it does not experience the damage mechanisms, such as “pitting”, intrinsic to solid targets. Finally, the source exhibits exceptional spot stability, with a spot centroid standard deviation of $< 0.1 \mu\text{m}$ over 24 hours:



Measurements of the Excillum MetalJet’s spot centroid over 24 hours, as taken with a pinhole camera mechanically coupled to the source. The standard deviation of position is $< 0.1 \mu\text{m}$ over 24 hours.

In the interest of transparency, we note one current boundary of MetalJet sources: the maximum beam voltage is 160 kV, while conventional X-ray sources typically run at 225 kV or even 450 kV. Beam voltage shapes the photon energy spectrum, with higher energies enabling penetration of larger and denser objects. For region-of-interest scans such as prismatic cell corners, and for full

scans of cylindrical cells including the larger 46XX series, 160 kV is more than sufficient. For high-throughput, full-cell scanning of very large prismatic and pouch formats, higher kV may be required. Excillum is actively investigating higher-voltage MetalJet sources to extend these benefits to even larger cell formats.

Long story short, the MetalJet is unquestionably the best X-ray source for high-throughput battery quality control. If you need to scan as many cells as you possibly can with high image quality, you'll almost certainly need a system with a MetalJet X-ray source.

Now, with an understanding of the superscanner's X-ray source, we turn the VCB2 itself.

The VCB2 “superscanner”

Excillum’s MetalJet X-ray source is unquestionably the breakthrough innovation enabling ultrafast cell CT scanning, but building a high-speed, production-grade CT system based on the MetalJet is a massive undertaking in its own right. For this task, [VisiConsult](#)’s deep expertise in developing custom CT systems with integrated automation solutions made them the perfect CT system integrator to bring the superscanner to life.



Left: Mockup of the VCB2 design. **Right:** The first VCB2 system.

We’ll begin by discussing the superscanner’s CT imaging chain.

Imaging chain: X-ray detector

An immediate challenge introduced by the MetalJet’s ultrahigh brightness is pairing it with an ultrafast X-ray detector to capture this signal. Most industrial CT systems sold today use “[indirect detectors](#)”, meaning that they convert X-rays to electronic signal by way of a scintillator. Scintillators convert X-rays to visible light, which is subsequently converted to electronic signal. Indirect detectors are used in essentially every industrial CT system on the market today. These detectors have a maximum frame rate of ~50 frames per second. In contrast, [direct detectors](#) directly convert X-rays to electronic signal, not unlike a solar cell. Direct detectors (a special flavor of which are “[photon-counting detectors](#)”) are often much faster (up to [1000 frames per second!](#)) and are thus a natural choice to pair with the MetalJet X-ray source.

For the superscanner, VisiConsult has worked closely with [Varex](#), the industry leader for X-ray detectors, to integrate their [DC-THOR](#) photon counting detector. This detector unlocks high-resolution (100 μm pixel pitch) acquisition at up to 500 frames per second in continuous acquisition mode. The

combination of the MetalJet X-ray source and the DC-THOR X-ray detector enables the ultrafast scan times enabled by the superscanner. To the best of our knowledge, the VCB2 is the first industrial CT scanner in the world to incorporate photon-counting detectors (although this technology has been used in medical CT for over twenty years).




The DC-THOR photon-counting detector, capable of continuously acquiring 500 frames per second.

Imaging chain: CT acquisition modes

The superscanner is primarily designed for helical acquisition.


Conventional (“cone beam”) CT scanning uses a circular acquisition trajectory, meaning the part rotates on the stage at fixed height. In contrast, during helical acquisition, the part rotates and translates on the stage simultaneously.



Left: “Cone-beam” (circular) acquisition. **Right:** Helical acquisition. ([Image source](#))

Helical acquisition offers two distinct advantages over conventional circular acquisition:

1. First, helical acquisition enables higher magnification than circular acquisition because the object moves continuously along the rotational axis, meaning we don't need to fit the entire object into every single 2D projection but can instead sample it continuously, which allows the object to be positioned closer to the X-ray source. Higher magnification translates to smaller voxel size.
2. Second, circular reconstruction algorithms typically introduce "cone beam artifacts" (see below; read more [here](#)). Unfortunately, for battery scans, these artifacts tend to interfere with critical regions of the cell such as the overhang region. While some measures can be taken to decrease the severity of these artifacts in circular acquisitions, like tilting the cell, helical acquisition & reconstruction can avoid these issues entirely.



Left: "Cone beam artifacts" in the reconstructed volume of a cylindrical cell scanned via circular acquisition. **Right:** Artifact-free helical reconstruction of the same region of the same cell scanned via helical acquisition.

VisiConsult has also developed advanced, automatic geometric calibration routines, as poor geometric calibration is a common source of [imaging artifacts](#), known as "[windmill](#)" artifacts, in helical CT scanning. The calibration method calculates all relevant geometrical reconstruction parameters and passes them directly into the reconstruction pipeline.

All that said, the superscanner also supports conventional circular ("cone beam") CT acquisition as well. For instance, customers may opt for circular scanning if they are only interested in small, targeted region(s) of interest such as the top and/or bottom "overhang" regions of a cell. Again, we generally recommend

full-cell scanning since defects and quality issues can occur anywhere in the cell, but region-of-interest scanning makes sense in some scenarios.

Imaging chain: Mechanical and thermal control

While we'll discuss part handling in depth momentarily, the mechanical and thermal control of a CT system can significantly impact its image quality.

VisiConsult has carefully considered these aspects:

- The system mechanics are carefully engineered to maintain subpixel accuracy ($\leq 20 \mu\text{m}$). Some of the key mechanical control features are a unique axis calibration algorithm and high-precision rotary encoders. The system is [VDA5](#)-compliant.
- The system has active thermal management to maintain ultrahigh temperature stability over time.


Imaging chain: Reconstruction

The final piece of VisiConsult's imaging chain is [reconstruction](#). As a reminder, a reconstruction algorithm converts the sequence of thousands of 2D images into a 3D volume. Reconstruction is a computationally-intense process, often requiring 1–20 minutes per reconstruction even with cutting-edge compute hardware such as the latest GPUs. VisiConsult's reconstruction software can reconstruct high-resolution, full-cell CT scans in <10 seconds, which is likely the fastest reconstruction time for any industrial CT scanner.

Reconstruction algorithms often contain several subroutines, each of which must execute properly to produce a proper reconstruction. While the stability of these algorithms is generally high, even 99% stability would result in 100 scans/day with reconstruction errors. Thus, VisiConsult has implemented multiple automated quality checks for flux, edges, and alignment to confirm only high-quality scans are outputted.

Sample scans


Some sample images from a 2170 scan are displayed below. This scan is a full-cell scan with $25 \mu\text{m}$ voxel size obtained in 10 seconds. At the time of this writing, VisiConsult has a few more short-term levers to pull that will further reduce scan time while maintaining high image quality.



Sample cross-sectional images from a 10-second, full-cell scan of a 2170 cylindrical cell with 25 μm voxel size. These scans can be viewed in full via "Demo Mode" in the Glimpse Portal ([link](#)).


Part handling

First, we note that due to the MetalJet's dual X-ray beam functionality, each superscanner has two imaging chains ("left" and "right"), leading to 2x throughput compared to a single-beam design. The X-ray source is in the center of the system, and two sample stages and two detectors are located on either side of the source.



Depiction of the VCB2's two imaging chains. Powered by the MetalJet's dual-beam functionality, the VCB2 has two sample stages and two detectors to maximize throughput. In this design, the MetalJet is in the center of the system, and the two imaging chains ("left" and "right") are on either side of the source.

Of course, fast scan time also requires fast cycle time, meaning the part handling "dead time" must be minimized. Conventional CT scanners are rarely built for rapid part handling, often requiring opening and closing the scanner door when (un)loading each cell. VisiConsult's design enables efficient part conveyance via its innovative rotary table design. Each of the two rotary tables has seven cell positions: two for cell (un)loading outside the scanner and five in the scanner. The beauty of this design is that it enables continuous operation of the scanner, since cells can be (unloaded) simultaneously to scanning. Furthermore, this turntable design enables easy integration with in-line automation systems. Lastly, cell barcode scanning is integrated into the cabinet and performed during CT scan acquisition.




Left: The VCB2's two rotary tables, each with seven positions. **Right:** The rotary tables as integrated into the VCB2 cabinet. When a cell is ready to be scanned, it rises from the rotary table into the path of the X-ray beam. The rotary tables design enables continuous operation of the scanner.

The superscanner's part handling capabilities enable a "dead time" of ~10 seconds per scan. Since the scan time ranges from 5–10 seconds, the cycle time is thus 15–20 seconds—per imaging chain! Thus, the system cycle time is 15–20 seconds / 2 imaging chains = 7.5–10 seconds per imaging chain, or 6–8 cells per minute. This cycle time is remarkable for high-resolution, full-cell CT scanning. VisiConsult is continuing to innovate on the superscanner's part handling time, refining the mechanics to minimize all possible "dead time", and expects to achieve 16 cells per minute in the next design.

Production considerations: Reliability, serviceability, and automation

The VCB2's intended use case in a production environment necessitates ultra-high reliability, and VisiConsult has risen to the challenge. Each component of the VCB2 has been selected for reliability, from the X-ray detectors to the precision rotary encoders. While long-term system testing is still underway, the initial reliability data indicates the system is working as intended.

Additionally, VisiConsult has designed the VCB2 for easy serviceability. For instance, the system has five service doors, including a dedicated door to access the tray-mounted MetalJet source:



Illustrations of the service tray and door for the MetalJet. The tray-mounted design enables easy servicing.

Finally, the VCB2 is designed to be integrated into in-line or at-line production. While the system is sold without automation capabilities (since each customer will have bespoke automation needs), VisiConsult will work with automation partners to develop seamless mechanical integrations.



Illustration of the VCB2 integrated into a production system. In this illustration, the scanner is integrated with a process step in which the cells are conveyed in trays, such as the formation step. A [“pick and place” robot](#) (un)loads cells from trays into the VCB2’s rotary tables.

Overall, the VCB2 combines cutting-edge X-ray imaging components with VisiConsult’s deep expertise in CT system design, automated part handling, and CT reconstruction. The result is a production-ready CT scanner capable of 10,000 high-resolution, full-cell scans of cylindrical cells per day.


The Glimpse Portal

We now turn to the final component of the superscanner solution: the [Glimpse Portal](#). We'll begin by discussing one of the core challenges of the superscanner: the truly massive volume of data produced.

A single high-resolution, full-cell CT scan of a cylindrical cell is around 50 GB. (One-sentence proof: a CT scan of a 4680 cell with a 25 μm voxel size at 32 bits/voxel will have a file size of $46\text{ mm} \times 46\text{ mm} \times 80\text{ mm} / (25\text{ }\mu\text{m})^3/\text{voxel} \times 4\text{ bytes/voxel} = 43\text{ GB}$, not including any padding to account for imperfect cell centering.) At six scans per minute, the superscanner would generate 300 GB/minute, 18 TB/hour, 432 TB/day, and 158 PB/year. Assuming [typical AWS storage costs](#) of \$0.02/GB/month, storing one month's worth of data would cost \$260,000/month. The compute and networking costs to handle this data would be considerable as well—not to mention the time to analyze this data. (Imagine the number of operators required to manually check for dozens of features in 10,000 scans each day...) Thus, processing, storing, managing, and analyzing this massive yet incredibly rich datastream is an essential challenge when operating the superscanner.

[Glimpse](#) was founded in 2023 to tackle exactly this challenge: enabling battery quality at scale by developing software solutions for high-throughput CT scanning. Glimpse's value proposition can be summarized as follows:

1. Glimpse's software improves the scan time–image quality tradeoff via its image enhancement algorithms.
2. Glimpse's software compresses the raw data with a compression ratio of up to 1000x so that this data can be stored in perpetuity.
3. Glimpse's software manages the data such that all scan data and metadata is organized and searchable in seconds.
4. Glimpse's "automated inspection" computer vision algorithms extract key defects and features from the CT scan images.
5. The [Glimpse Portal](#)[®] provides an intuitive interface for visualizing CT scan images and their associated automated inspection algorithm results. Furthermore, this web-based application enables sharing and collaborating on CT scan data across all authorized users within an organization.

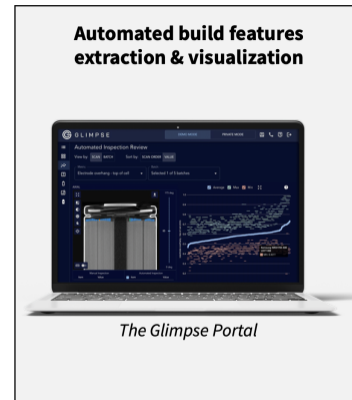


Pain-free scan data management

Scan images and metrics available:

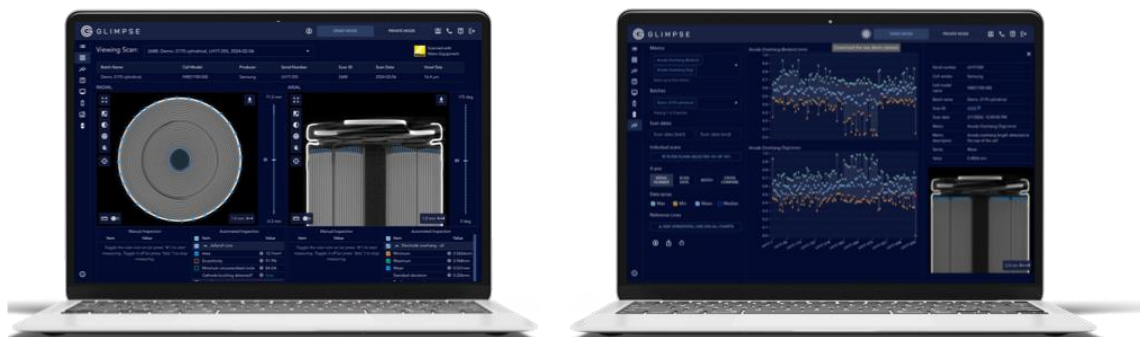
- In a few clicks,
- From any authorized company laptop,
- Anywhere in the world

Enabled by Glimpse's compression



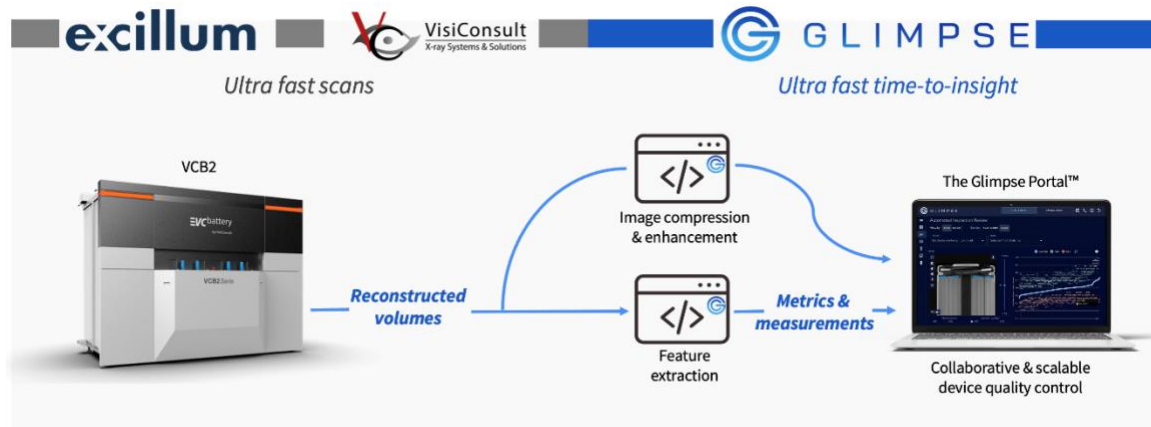
Summary of Glimpse’s value proposition. Glimpse improves the [Pareto front](#) of scan time and image quality, enables pain-free scan data management (enabled by Glimpse’s powerful compression algorithms), enables automated extraction of build features and defects, and provides an easy-to-use interface for visualizing this data (the [Glimpse Portal](#)).

The result of these innovations is a solution that minimizes the “time-to-insights” from the massive volume of data generated by the superscanner. Specifically, the Glimpse Portal is designed such that trends and outliers in thousands of CT scans can be identified in seconds. Check out Glimpse’s [free 1000-cell demo](#) of the Glimpse Portal to learn more. The demo includes some sample 10-second 2170 scans from the superscanner.



Screenshots of the Glimpse Portal. **Left:** Screenshot of the “Scan Viewer” page, designed to view a single scan along with the results of its automated inspection algorithms. **Right:** Screenshot of the “Automated Inspection Review” page, designed to view the automated inspection results of thousands of cells at once and quickly identify outliers and trends. A free demo of the Glimpse Portal can be found at <https://app.glimp.se>.

For the superscanner, VisiConsult and Glimpse have worked closely together to ensure a seamless data handoff from VisiConsult's reconstruction to Glimpse's scan processing pipeline. Here, Glimpse handles image enhancement, image compression, data management & visualization, and feature & defect extraction.



Simplified diagram of the data processing flow. The VCB2 produces reconstructed scans from its ultra fast scans. This dataset is then passed to Glimpse's pipeline, which performs both image compression and enhancement as well as feature extraction. Both the compressed/enhanced image stack as well as the results of the feature extraction are then passed to Glimpse's backend, and the results are available for either human consumption (via the Glimpse Portal) or machine consumption (via Glimpse's APIs).

Of course, the automated inspection computer vision algorithms are an essential component of the full superscanner solution. In an R&D setting, manual scan review via dedicated desktop software such as Volume Graphics is a perfectly suitable approach; at 10,000 scans per day, manual review becomes completely infeasible. Today, Glimpse builds and maintains automated inspection algorithms targeting our customers' key quality features, such as core circularity and anode overhang. Our typical development workflow is to align on a customer's defects and features of interest, identify the scan time required to accurately capture these defects and features, and then build algorithms with images from this agreed-upon scan time.

A few key points regarding Glimpse's "automated inspection" algorithms:

1. We often use both classical and AI-based [computer vision](#) approaches, depending on the task and customer requirements. Unlike perhaps other startups leveraging AI, we take data security very seriously; customers own their data, period.
2. While many providers offer algorithms for commonly-requested algorithms such as anode overhang, these approaches are very often locally-developed, hand-tuned algorithms without dataset versioning, model versioning, unit and end-to-end testing, performance monitoring, etc. As discussed in Google's seminal paper "Hidden Technical Debt in Machine Learning Systems" (⁷), algorithms developed without this supporting infrastructure often perform well in demo evaluations but fail in real-world production environments. We consider Glimpse's MLOps infrastructure to be "the machine that builds the machine", or the system that enables us to build, deploy, and maintain dozens of algorithms in a high-reliability production environment. Our experience has taught us that this infrastructure is simply nonnegotiable for reliable insight extraction at superscanner scale.

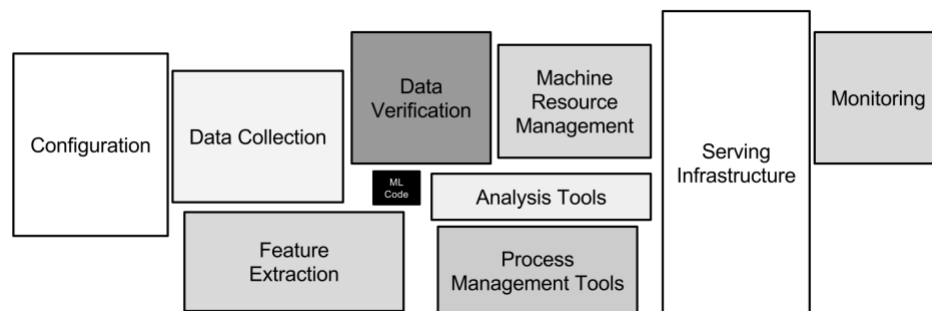
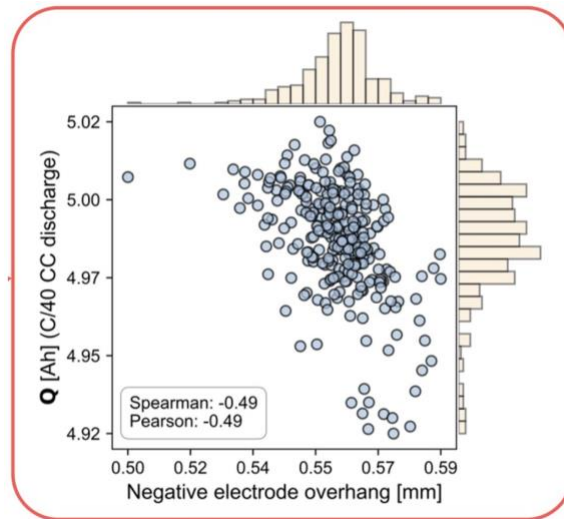


Figure 1 from Sculley et al.⁷, illustrating the scale and complexity of required supporting infrastructure when developing real-world ML/AI systems. Only a small fraction of real-world ML systems is composed of the ML code.

3. Obviously, these automated inspection algorithms are not very useful if their inference time cannot keep up with the rate of data generation from the superscanner. The minimum analysis time required is the time between scan completion and the cell exiting the system; assuming a 5-second scan time and thus a 15-second cycle time per imaging chain, this minimum inference time is 5×15 seconds = 75 seconds since a cell will go through five positions on the turntable before the cell exits the system. This analysis time includes reconstruction, image processing, and automated inspection algorithm inference. Glimpse will design a compute system to meet a customer's analysis time requirements depending on the number of algorithms, the complexity of the algorithms, and the customer's IT requirements (i.e., local vs. cloud compute, networking availability, etc.).

Finally, we return to our previous discussion of why this system is not designed to be simply a pass/fail solution. As we've discussed, despite the superscanner's breakthrough speed, its throughput is still much too low for 100% testing for large battery production facilities. Performing pass/fail testing on a single-digit percentage of cells (or less) is not particularly effective at reducing the rate of defective cells that escape into the field. Instead, the superscanner coupled with the Glimpse Portal unlocks continuous process improvements via an approach combining CT sampling with data analysis. By intelligently sampling cells to scan from across production, defects and concerning trends observed via CT (e.g., an increase in anode overhang violations) can be root caused to specific production steps and/or lines; similarly, anomalies in cell electrochemical performance (e.g., a decrease in cell energy) can be root caused by correlating to geometrical features observed via CT (see below). To this end, Glimpse provides APIs for both importing production data into the Glimpse platform and exporting data from Glimpse into a production data management system (e.g., a [MES](#)).



Correlation between cell capacity and negative electrode overhang. Obtained from a publication (in review) from Glimpse and a leading academic battery lab.

Conclusion

The VCB2 “superscanner” is the world’s fastest industrial CT scanner. Powered by the unique capabilities of [Excillum](#), [VisiConsult](#), and [Glimpse](#), this system is capable high-resolution, full-cell cylindrical cell scans in 5–10 seconds per cell. But this groundbreaking scan time is only half of the story: the VCB2’s ultrafast reconstruction coupled with the [Glimpse Portal](#) enable rapid “time-to-insights” from the massive quantity of data generated by this system. Furthermore, this system is designed for automation capability, readability, and serviceability to meet the exacting demands of production environments.



Summary of the key functionalities, components, and companies behind the superscanner.

After years of hard work, the superscanner is now a real, live, production-grade system. If you’re interested in learning more:

- Attend our webinars. We are hosting two webinars explaining both [technical](#) and [commercial](#) aspects of this system. The live webinars will be recorded in case you miss them.
- Reach out to us to learn more. We’d love to hear from you whether you’re interested in the VCB2, a similar system for another high-throughput QC application, or simply wish to learn more.
- If you have serious interest in this system, please reach out to coordinate a visit to VisiConsult’s HQ to see this scanner in action. Please note that we expect the first unit to ship to a customer in November 2025, so visits are only available until the end of October 2025. As an added incentive to

visit, VisiConsult is located near the beautiful medieval German town of [Lübeck](#).

We hope this white paper has been both engaging and informative, showing why we believe the superscanner will transform battery quality control.

References

1. Attia, P. M., Moch, E. & Herring, P. K. Challenges and opportunities for high-quality battery production at scale. *Nat. Commun.* **16**, 611 (2025).
2. Min, J., Condon, A. & Attia, P. M. Lab-Scale X-Ray Exposure Has No Measurable Impact on Lithium-Ion Battery Performance and Lifetime. *Batteries* **11**, 73 (2025).
3. Condon, A. et al. A dataset of over one thousand computed tomography scans of battery cells. *Data Brief* **55**, 110614 (2024).
4. Hemberg, O., Otendal, M. & Hertz, H. M. Liquid-metal-jet anode electron-impact x-ray source. *Appl. Phys. Lett.* **83**, 1483–1485 (2003).
5. Danielsson, M., Persson, M. & Sjölin, M. Photon-counting x-ray detectors for CT. *Phys. Med. Biol.* **66**, 03TR01 (2021).
6. Davis, R. First Big U.S. EV-Battery Plant Offers Lessons as Industry Springs Up. *Wall Street Journal* (2023).
7. Sculley, D. et al. Hidden Technical Debt in Machine Learning Systems. in *Advances in Neural Information Processing Systems 28* (2015).