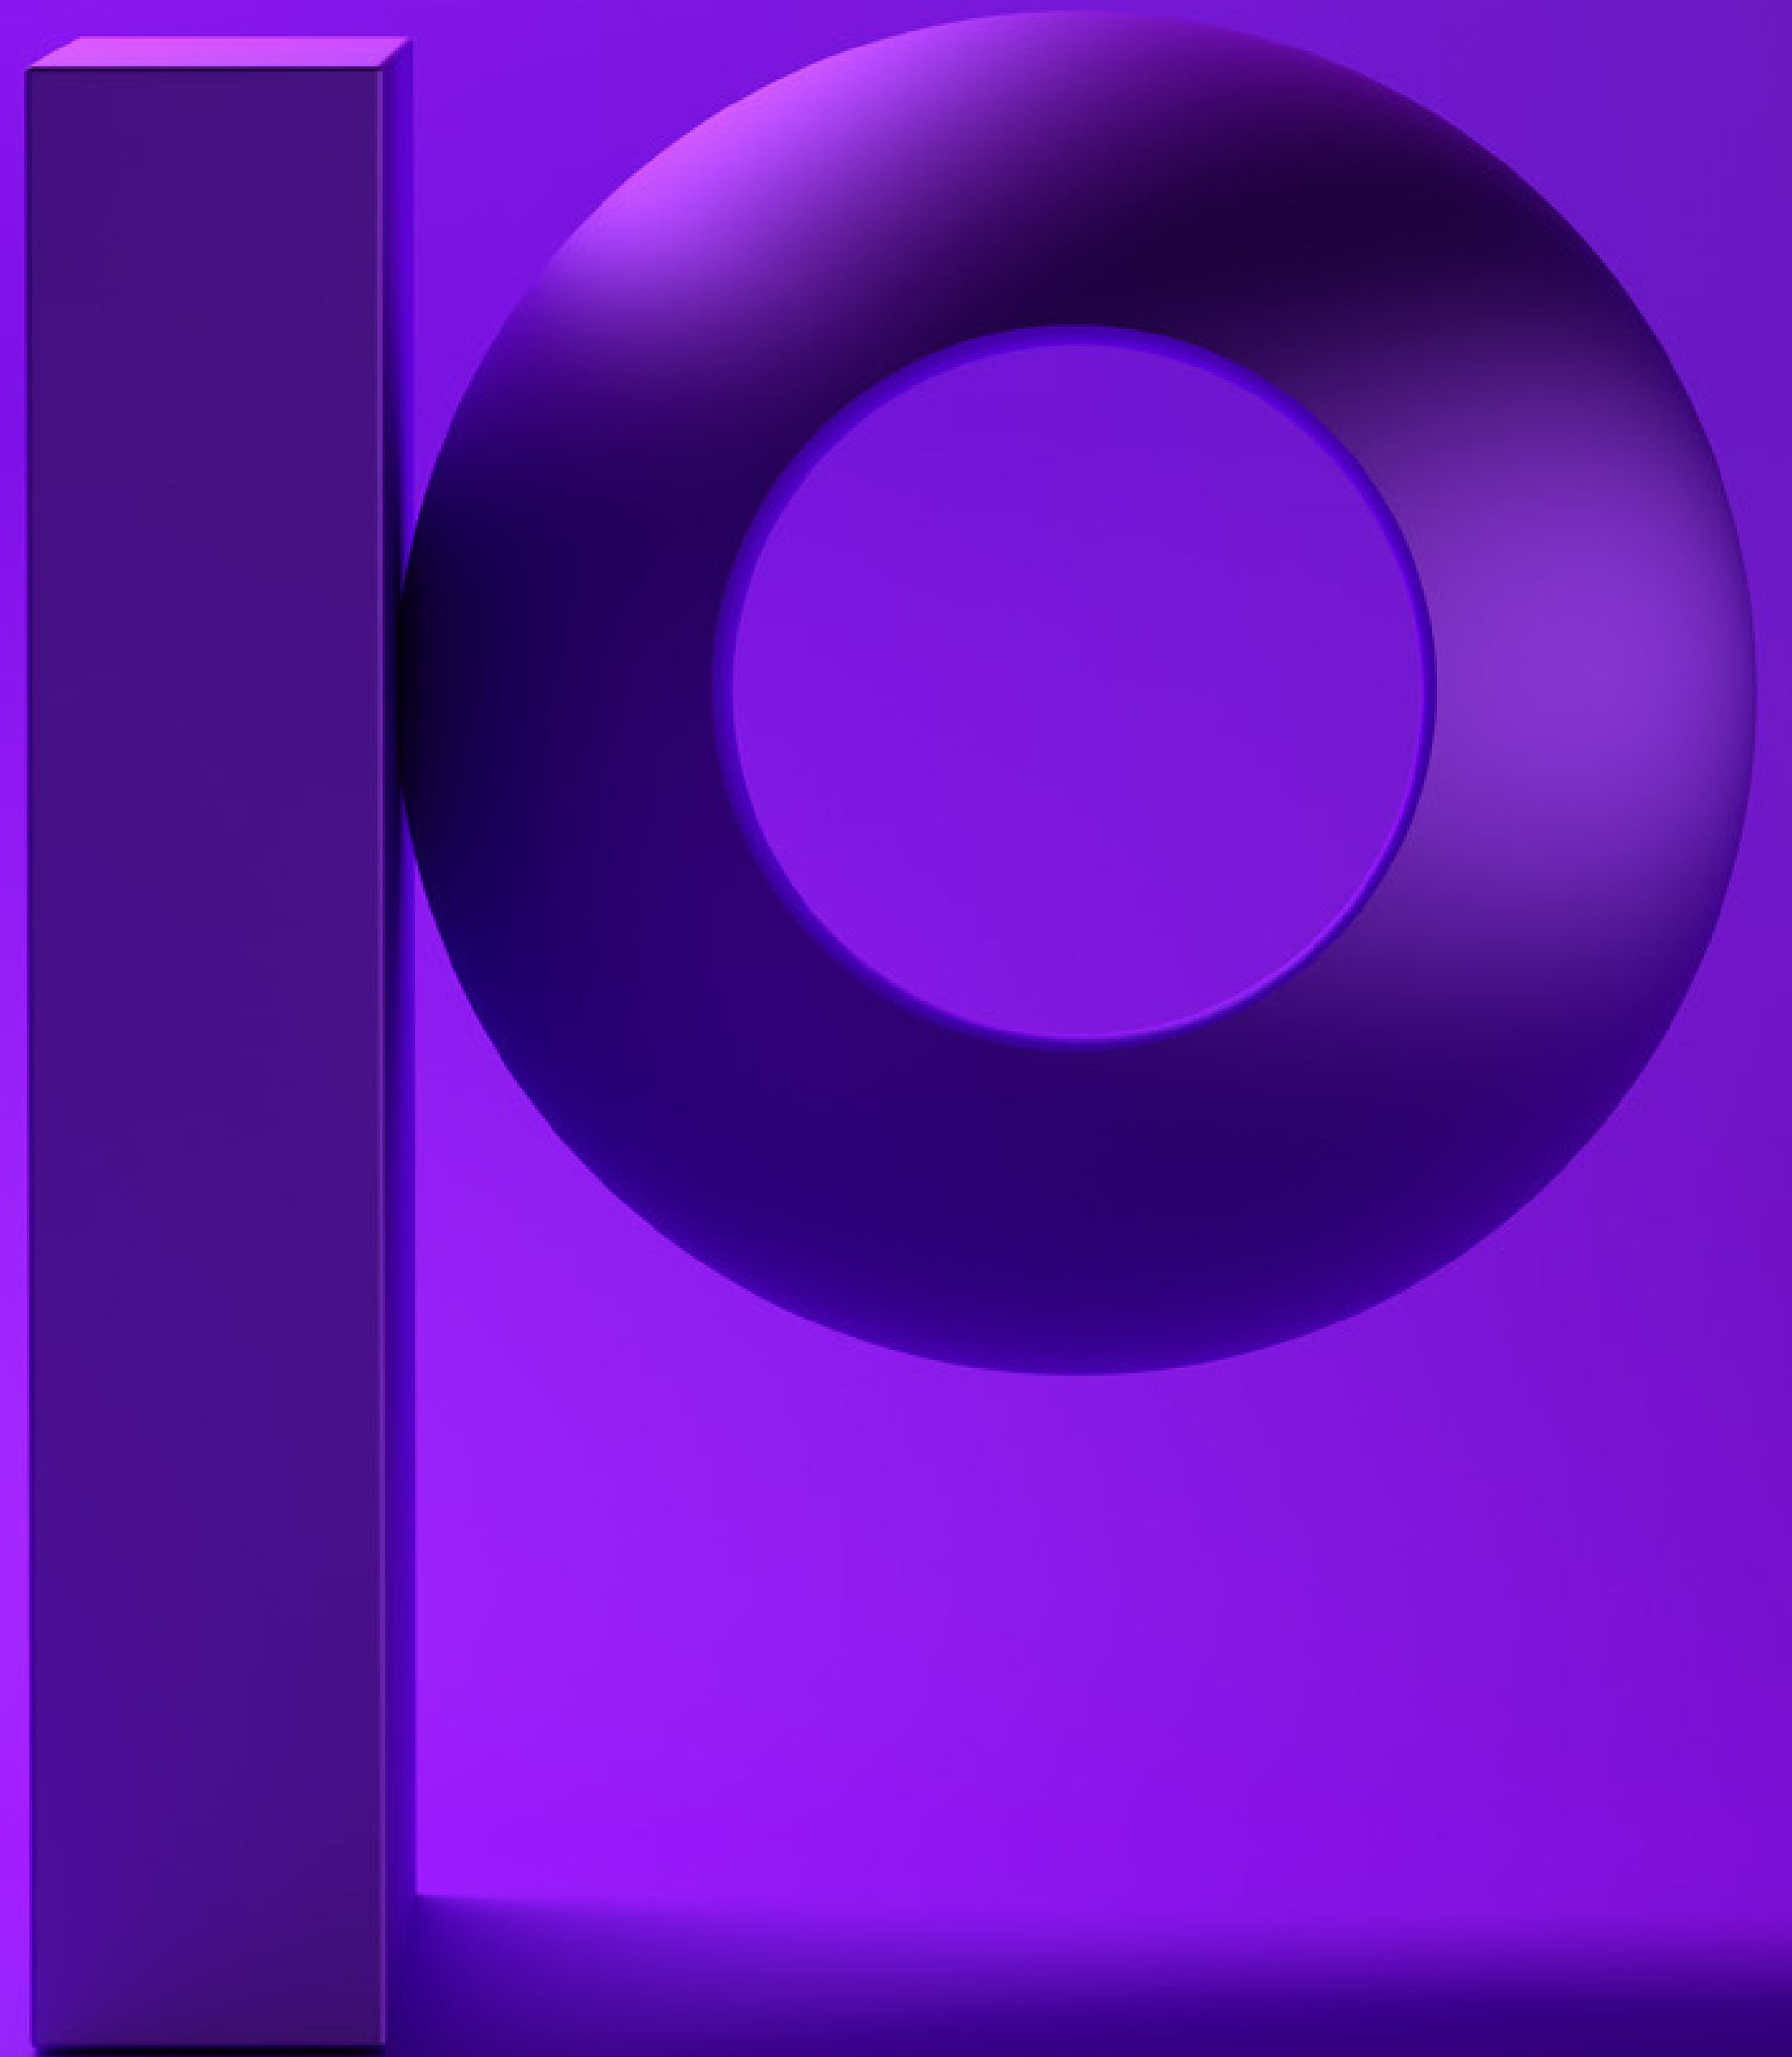precisely

Next-Generation Network Planning

The critical role of raster data in fueling the telco networks of tomorrow

Network planning at 5G speed

It was dubbed '4G Long Term Evolution' for good reason. From the start, it was clear that fully implementing 4G communication systems was going to take time for many reasons, including the fact that the available equipment and computing technologies were not yet capable of fully delivering on the promise.

In comparison, 5G is more like "The Big Bang." 5G network deployment and end-user device upgrades are happening with unprecedented speed. Intense competition is pressurizing the race to secure the best tower sites and most profitable customers. Any delays, any errors in deployment factors such as siting or modeling for Radio Frequency (RF) propagation can quickly result in escalating costs as well as lost revenues and customers.

The challenges of 5G network deployment are multi-faceted and multi-layered, both figuratively and literally. Just the sheer number of technical factors that need to be considered can be mind numbing. For example, 5G's extremely short mmWave signals have a much shorter effective range, are more susceptible to multipath distortions and local interference, and are also more severely attenuated by buildings, landscapes, and weather conditions. While technologies such as beamforming and MIMO transmission can surmount these problems, deploying them properly requires encoding and analyzing a much greater volume and granularity of spatial data points and details.

Network planning involves much more than engineering the deployment of the physical technology. The resulting network must be a robust and extensible platform for enabling current and future service offerings, which means planning and deployment are not one-time efforts. It is an ongoing, constant chase to identify and respond to infrastructure changes on the ground, network capacity requirements, and your customers' ever-shifting service and Quality of Service demands.

Ultimately, successful network planning and deployment depend on your ability to analyze huge volumes of raster file-based GIS data, in both spatial and temporal contexts. Raster files really are the fuel your network planning systems run on. Yet it can feel like that fuel is a thick, heavy sludge when you work with it. In addition to the sheer volume of data that must be processed, the broad diversity of file formats, granularity, and spatial extents in which the data are delivered makes analysis complicated and frustratingly compute-intensive.

The usual remedy for these challenges has been ever-increasing spending on more computing capacity. But given the explosive growth in raster data volumes required for 5G network planning, this IT budget-busting solution is unsustainable. To keep up with the demands of 5G network deployment, the processes for raster file data analysis need to undergo a transformational shift that is on par with the revolutionary advances in 5G technology they are supporting.

The data drag on 5G network deployments

5G network planning involves balancing and optimizing a bewildering number of business and technical requirements and constraints. Except for a few "greenfield" deployments, the core exercise of 5G deployments is heavily weighted toward the in-fill of existing 3G and 4G deployments with a 5G small cell "mesh."

Achieving that requires analyzing extensive and detailed GIS data covering multiple layers of topology and civil infrastructure as well as location information for populations, businesses and network traffic, 3D details of building construction, and more, including:

• Existing legacy network and system deployments, including those of competitors
• Existing radio signal environments
• Competition for cell tower sites
• CapEx, Cost/Benefit of network infrastructure options

Most of the data that feeds network modeling systems is raster file-based. But despite that commonality, viewed as a whole, the universe of raster data is essentially unstructured and rather difficult to correlate and analyze. The barriers to raster file analysis include:

- The diversity of raster file formats – over 100 different file formats
- Dissimilar data definitions and table structures, tiling, null/empty cell presentation, etc.
- Wide variations in spatial footprints and extents, resolution/ granularity
- Complex, multi-layer data representation for spectral bands, time series, etc.
- Data packaging/presentation: File sizes and compression methods

The impact of all these factors is most clearly seen the moment you push the enter key to launch an analytics or GIS modeling run. The compute loads required for harmonizing, correlating, layering, and analyzing such volumes of essentially unstructured data can be staggering.

Just the pre-processing steps required to load and restructure gigabytes (and even terabytes) of data into usable form for your analytics platform can take hours. And given the reality that you rely upon more than one analytics and visual presentation platform, there will also be a heavy processing load and latency involved as you manipulate and model against the merged data. Meanwhile, as the deployment of 5G continues to ramp up, the raster file processing overload issue will only get worse.

Geospatial Raster Data

Raster data is, in general, any pixelated (or 2D gridded) data, such as a photo. But geospatial raster files must also include spatial information that connects the data to a particular real-world location, including its extent and a resolution.

- The extent is the geographical area covered by a raster
- The resolution is the area covered by each pixel of a raster

Multi-band Raster Data
A geospatial raster file can contain one or more bands. An RGB color image is a basic three-band raster. But geospatial raster can contain many additional bands:

- **Time series:** record of the same data point or location, over time
- **Multi- or Hyper-spectral imagery:** image rasters that have 4 or more (multi-spectral) or more than 10-15 (hyperspectral) bands. These bands may include variables regarding weather conditions, radio-spectrum information, or statistics regarding call or data volumes

Multi-band raster data in telecommunications
Telecommunications companies continually process a huge number of call records. These can be classified and presented via highly accurate, near real-time coverage maps. This type of visual analysis helps firms to improve customer service, reduce churn, market more effectively, and gain market share.

Rethinking raster data for 5G network planning

Trying to resolve the issue through ever-expanding compute capacity is a fool's errand. Instead, the central problem, the incompatibilities across the universe of raster data files, must be overcome.

Breaking that problem down to its more fundamental elements, we can describe the key requirements and goals for raster file reform:

- Unifying/standardizing the presentation of diverse raster data types
- Simplifying the transformation of data resolutions and extents for efficient layering
- Significantly reducing data file volumes for improved compute efficiency and reduced data storage resource needs

precisely.com | 877 700 0970
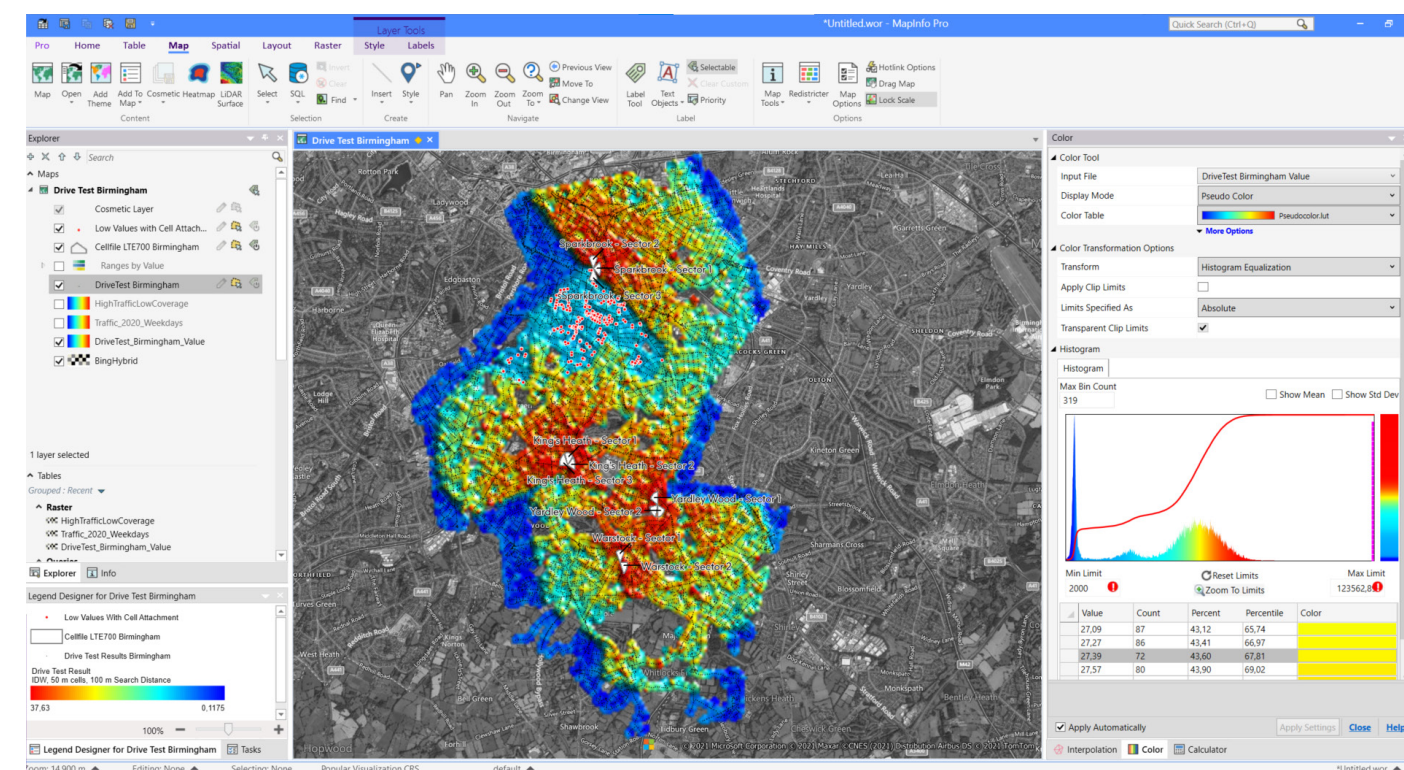
Unified presentation

There are currently more than a hundred raster data formats in use, embodying imagery, continuous gridded data, and thematic data. In addition, raster data may be expressed in bands and/or color pallets, or in numeric values, text strings, and more.

An increasing number of planning factors also require analysis across time. Consider call and data volumes, equipment replacement, and upgrade cycles, etc. Storing and presenting what are essentially multiple time-stamped versions of a base raster file rapidly inflates the dataset. Most often, a painful choice must be made between full temporal resolution (capturing data in smaller, more granular time periods) and the impact on compute loads and processing time.

What is needed is the ability to bring data of multiple types and formats, including temporal data, into a unified presentation format that can be used across applications and analytics methods.

Simplified transformation and scaling

Ignoring for a moment the problem of file format incompatibilities, the resolution or granularity of data stored in raster files varies widely. Likewise, the extents covered in individual raster files using the same format also varies widely, even across successive editions of the same dataset.

These variations make it difficult to align and harmonize multiple datasets into compatible layers for analysis. Manipulating the data for visual rendering at different scales or zoom levels also ends up requiring complex and time-consuming computation.

Here the need is to be able to easily and efficiently port raster data from its original resolution and extent into a more flexible, interoperable, and scalable layer for presentation and analysis.

Data volume reduction

The increased equipment density and number of planning factors involved in 5G networks inevitably expand the already huge volume of data which is required to perform sufficiently detailed analysis.

Raster file size depends upon the spatial extent being captured and the density of the data points within that footprint. As 5G networks expand their coverage, they do so with a higher density of equipment and so require many more data points to be captured in any given raster dataset.

IT resource limitations make it increasingly impractical to work with the resulting monolithic, high-resolution raster files. This has forced planners to accept a trade-off between raster resolution and spatial extent. The choice boils down to increased compute resources and run times or generating less detailed, less accurate results.

The preferred solution for this is to design ways to encode raster data more efficiently and intelligently, reducing, or eliminating brute-force read/write cycles across raster file tables and reducing their required storage volumes.

Multi-Resolution Raster: accelerating spatial data processing

To address these issues and enable more efficient and flexible raster data processing for 5G network planning, Precisely developed an innovative raster format called "Multi-Resolution Raster" (MRR).

Ultimately, storing and processing raster data in MRR format virtually eliminates the previously unavoidable inefficiencies and barriers to working with disparate and complex forms of raster data for next-generation 5G network planning, deployment, and operations management.
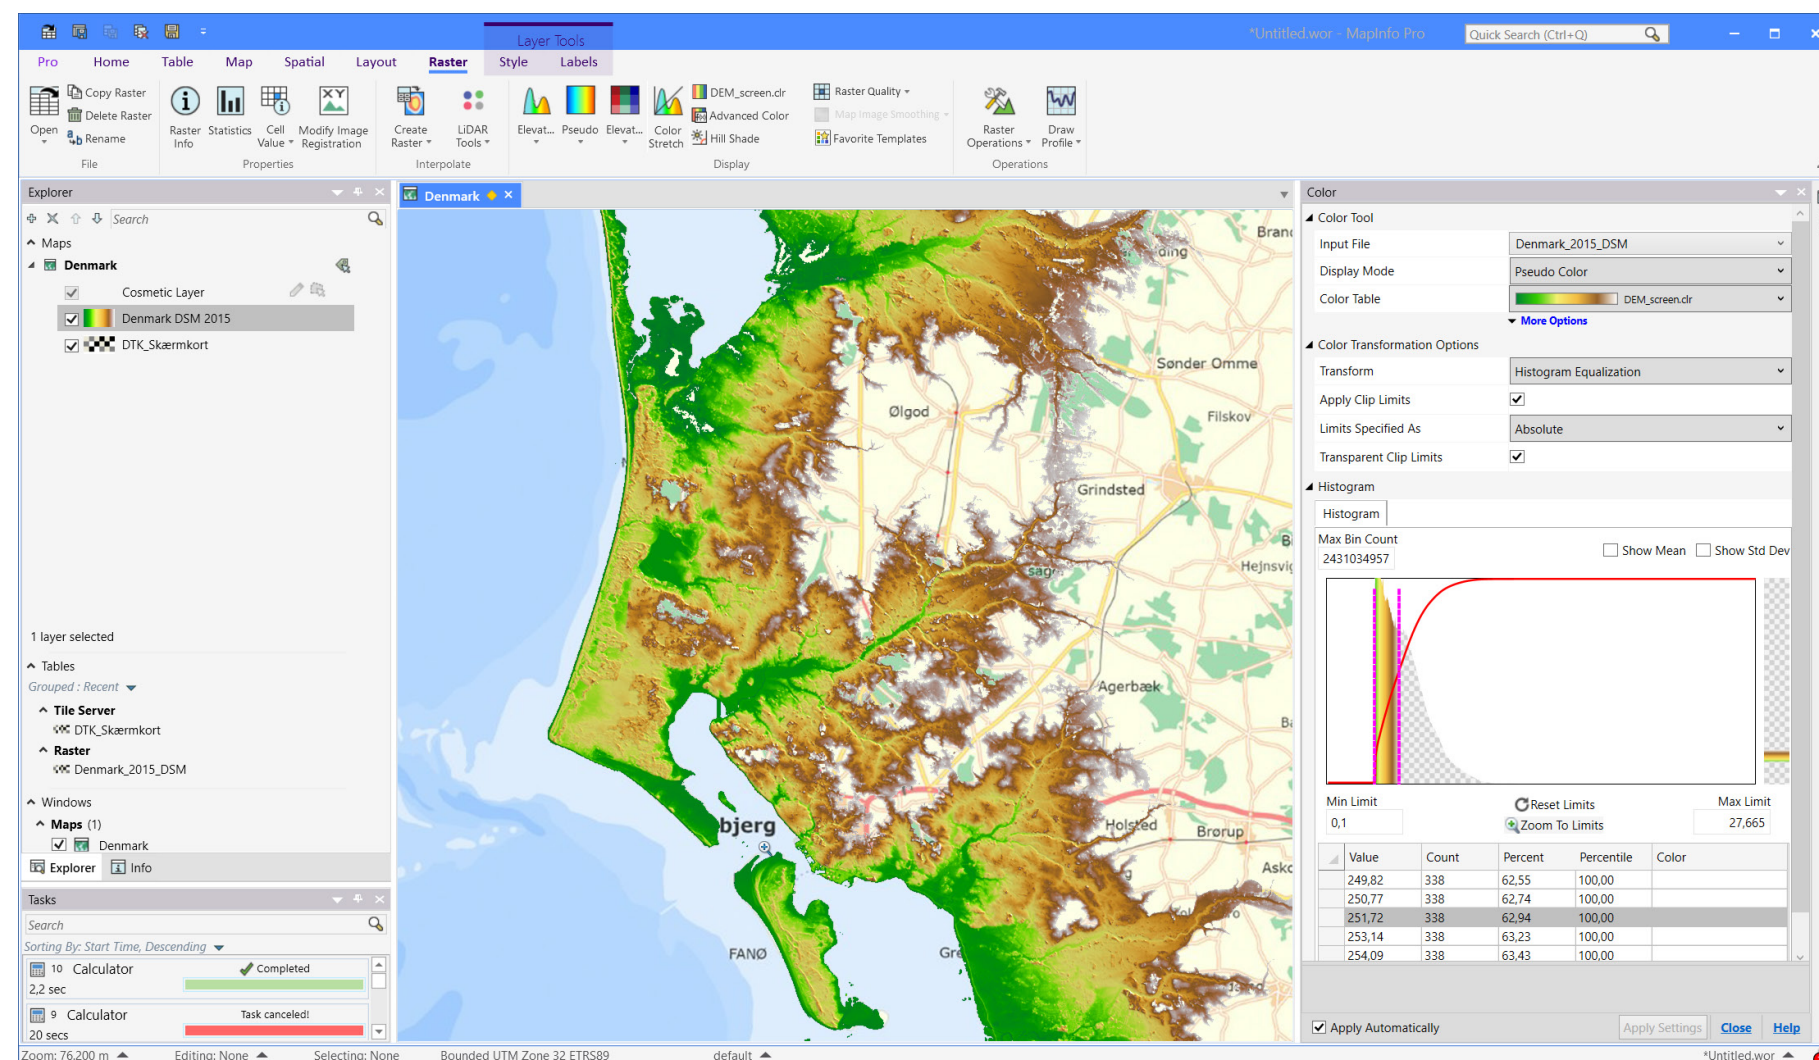
Unified presentation

With MRR, Precisely redefined not just how raster data is stored, it opened up the options for what kinds of network planning data can be stored in a raster. It was developed specifically to enable combined, unified storage of all forms of raster data, including imagery, spectral imagery, continuous gridded data, and classified data. Existing raster data files in any format can be easily combined into a single MRR file, without restrictions to raster file spatial extent, resolution, or complexity.

Simplified transformation and scaling

The MRR file format is designed to deliver an enhanced visualization experience by maintaining an overview data pyramid, making opening, panning, and zooming of MRR files near-instantaneous.

Every raster has a base resolution level from which the source data is accessible. But MRR files also store a multi-level pyramid of overview levels, at half resolution increments from the base. This allows applications fast, direct access to data from throughout the pyramid efficiently and flexibly, at whatever resolution is most efficient for the process.

In addition, MRR files also provide "under-views." In an under-view, data is presented at a higher resolution than the original data. Using a variety of interpolation techniques such as nearest neighbor, bilinear, and bicubic under-views are generated on the fly, at whatever resolution is required.

Taken together, the overviews and under-views provided within MRR formatted raster data files enable extremely fast and responsive, high quality visualization at any scale, with extremely low compute demands.

Data volume reduction

Efficient storage is a pre-requisite enabling technology for large raster file analysis and presentation. Physically, MRR files are stored on disk as a single unified file rather than as a set of MRR formatted files. But rather than creating an incredibly large, monolithic file, MRR actually reduces the total bytes required for storage.

In an MRR formatted file, each component raster file's data is not stored as a simple, physical 2D array of cells but rather as an extensible sparse matrix of tiles of cells. This sparse tile matrix arrangement drastically reduces overall storage requirements without losing any data or impacting its resolution. The combined effect of the sparse matrix storage method and the greatly reduced total file size radically simplifies and accelerates read/write access, resulting in exceptional gains in compute efficiency and speed.

MRR formatted files achieve the flexibility required to incorporate multiple large and disparate raster datasets into harmonized, compatible layers and bands, while preserving the ability to apply future edits and additions to each distinct dataset, but without creating an impossibly large data store that chokes network planning and modeling systems.

Case in Point: Building a trillion-cell MRR raster

The Danish government recently undertook two major LiDAR surveys over their entire country to enable creation of a high-resolution digital terrain model (DTM), capturing geologic topology, and a Digital Surface Model, a far more detailed record (sub-millimeter vertical resolution) which also included anything standing on the surface, from trees and buildings down to highway medians, cars, and even farm animals.

For the DTM, the resulting LIDAR image data totaled 668 ZIP archives, each containing up to 100 TIFF rasters. Each TIFF file captured about 1 square kilometer, resulting in well over 50,000 separate rasters. The resulting total storage requirement was 548 GB.

To reduce storage requirements and enable faster and more streamlined spatial analysis, the many thousands of TIFF files were merged into one MRR file and various MRR-supported techniques for compression were applied, including:

- LZMA level 9 data compression
- Scaling and translation of 4-byte floating-point data to 2-byte signed integers
- Perfectly reversible predictive encoding to enable data stream compression

The final MRR file, including the overview pyramid, was 164 GB, a 70% reduction.

The more detailed DSM dataset was larger, with 673 ZIP files, totaling 592 GB. After similar conversion to MRR format, the resulting DSM dataset was 226 GB, a net 62% reduction.
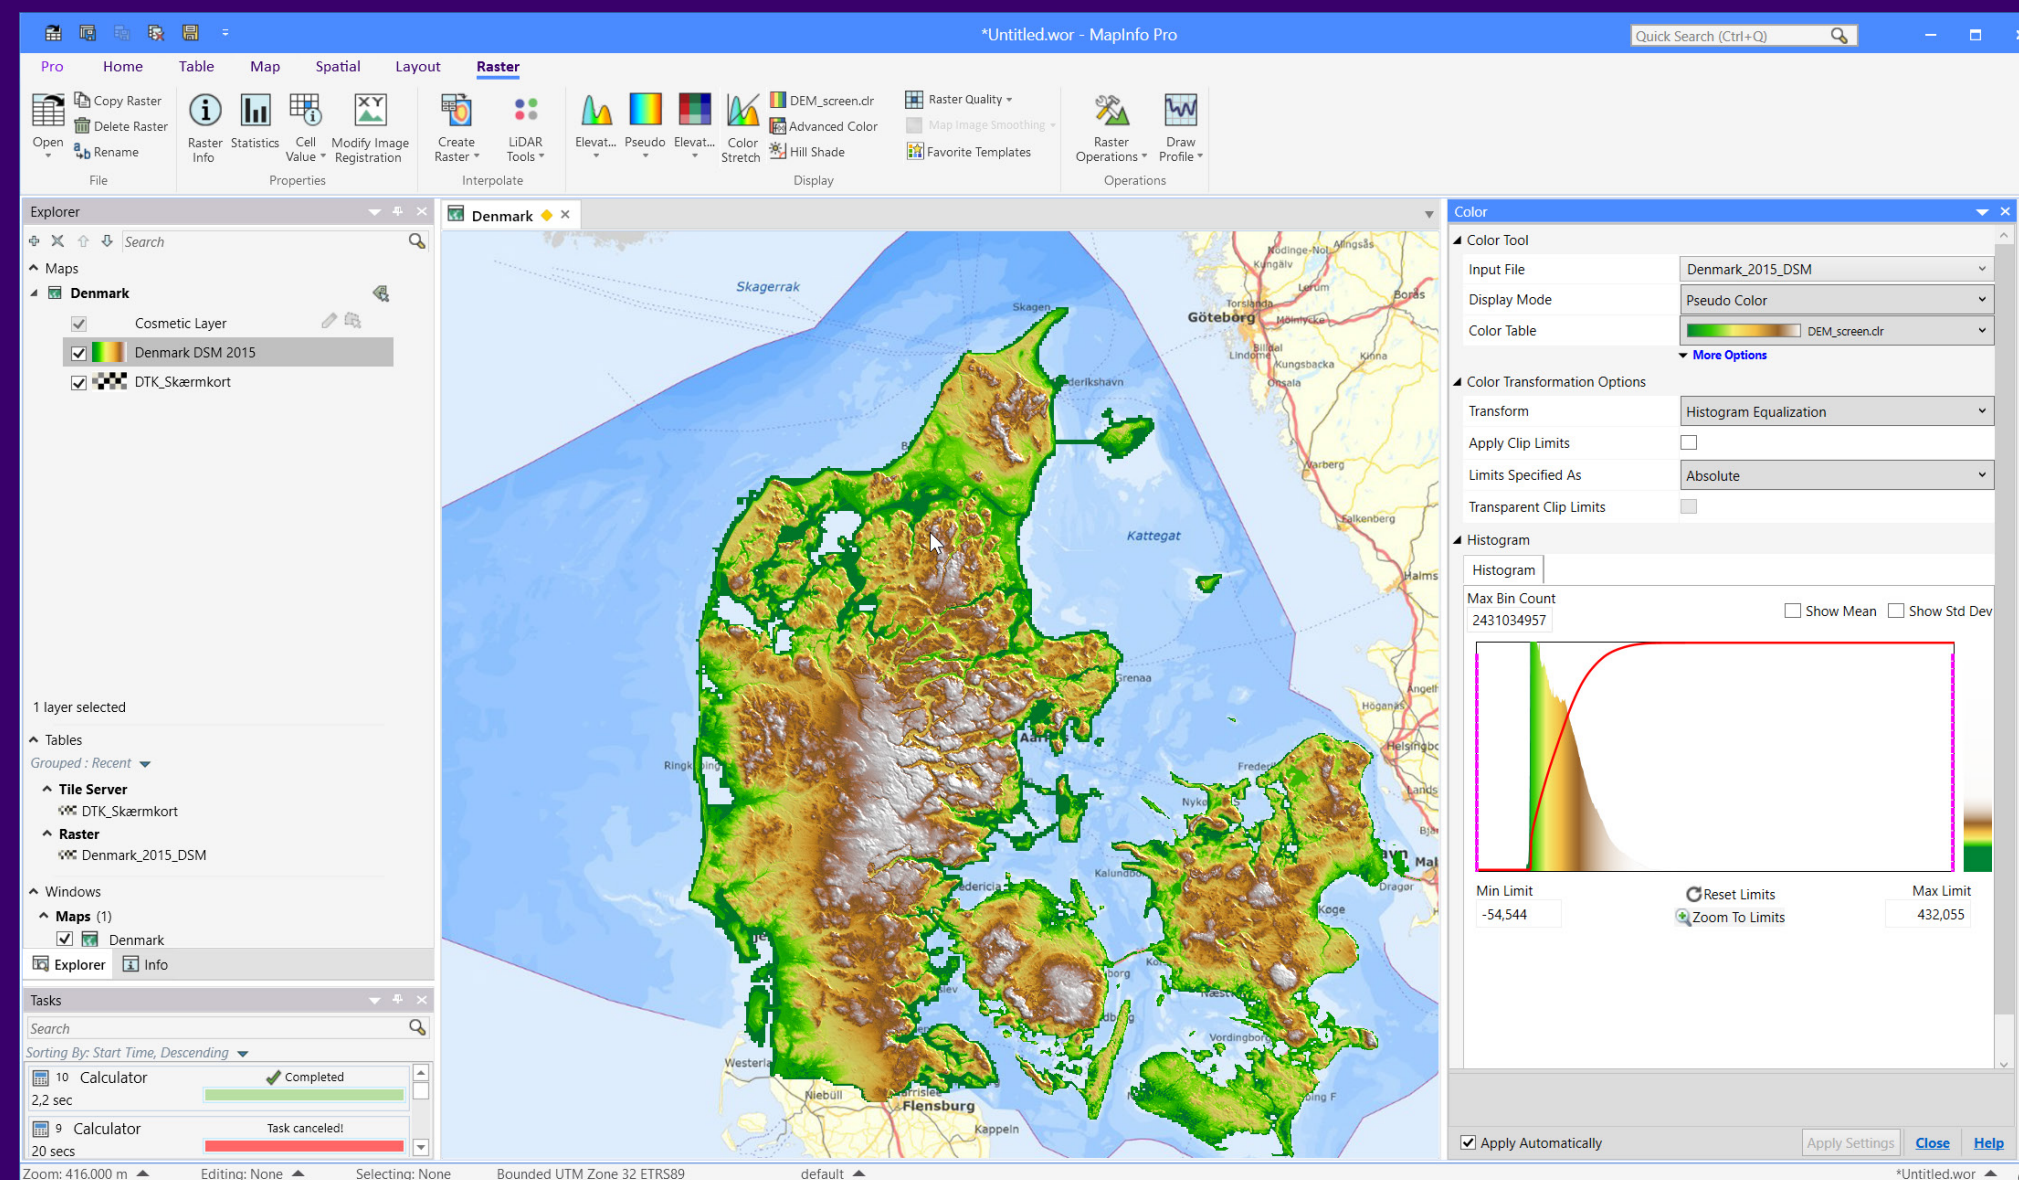
Next-gen networks call for a next-gen approach to raster data

5G communications technology will undoubtedly continue to change and improve how we all live, work, and conduct business for years to come. And for Communications Services Providers in particular, 5G technologies will open up a nearly endless stream of new business opportunities, and major headaches as well.

Just the increased density of network gear required and the accompanying shift to more complex and distributed control systems make network planning and deployment far more challenging than ever.

Moreover, the speed and reliability of 5G technology will drive rapid expansion of machine data communication services and mobile content delivery as well, which in turn will require savvy, carefully allocated network slicing and heavy reliance upon efficient and reliable overlay of 5G communications onto existing 4G core infrastructure.

Hitting the practical and physical limits of raster file data analysis and storage just as you are taking on these challenges is a serious and costly problem. Adopting Precisely's Multi-Resolution Raster file format is the most efficient and cost-effective way to ensure your systems are ready to meet the demands of next-generation network planning.

To learn more about why Precisely is a trusted provider of location intelligence solutions and data to major communication services companies around the world, visit our website. Or to try the raster functionality out for free, visit the MapInfo Pro webpage

precisely

About Precisely

Precisely is the global leader in data integrity, providing accuracy and consistency in data for 12,000 customers in more than 100 countries, including 90 percent of the Fortune 100. Precisely's data integration, data quality, location intelligence, and data enrichment products power better business decisions to create better outcomes. Learn more at www.precisely.com.

www.precisely.com